working paper
department
of economics

TACIT COLLUSION IN A DYNAMIC DUOPOLY WITH
INDIVISIBLE PRODUCTION AND CUMULATIVE
CAPACITY CONSTRAINTS

Glenn Loury

No. 557   May 1990

massachusetts
institute of
technology

50 memorial drive
cambridge, mass. 02139

TACIT COLLUSION IN A DYNAMIC DUOPOLY WITH
INDIVISIBLE PRODUCTION AND CUMULATIVE
CAPACITY CONSTRAINTS

Glenn Loury

No. 5561   May 1990

Tacit Collusion in A Dynamic Duopoly with Indivisible Production and Cumulative Capacity Constraints

by

Glenn C. Loury*

J. F. Kennedy School of Government
Harvard University**

June 1990

Abstract

This paper studies a dynamic, quantity setting duopoly game characterized as follows: Each firm produces an indivisible output over a potentially infinite horizon, facing the constraint that its cumulative production cannot exceed an initially given bound. The environment is otherwise stationary; the remaining productive capacities of the firms at any moment are common knowledge; the firms choose production plans contingent on these capacities which are mutual best responses in every contingency. The resulting Markov Perfect Equilibria are analyzed using a two–dimensional backward induction, and compared with the equilibria which emerge when precommitment to time paths of output is possible. It is shown that the ability to precommit can be disadvantageous; that collusion in Markov Equilibrium is facilitated by the symmetrical placement of the firms; and that having greater capacity confers basic strategic advantage on a firm by enabling it to credibly threaten future production. The model solves an open problem in the theory of exhaustible resource economics by imposing subgame perfection in a resource oligopoly with independent stocks. It also formalizes the intuition that, when indivisibilities are important, tacit coordination of plans so as to avoid destructive competition is facilitated by establishing a convention of "taking turns" — that is, a self–enforcing norm of mutual, alternate forebearance.

I. Introduction

This paper studies tacit collusion in a dynamic model of duopoly. Firms face an indivisibility in production at any moment, and an intertemporal constraint on the cumulative amount of output which is possible over the horizon of the game. The analysis explores the strategic implications of these two conditions: On the one hand, due to the indivisibility, a firm deciding whether to produce at a point in time has limited flexibility to accomodate its opponent's action. On the other hand, due (say) to the presence of an essential, limited and non–reproducible factor of production, each firm's capacity for future production is progressively diminished by decisions it has taken to produce the past.

We study the Markov Perfect Equilibria of the resulting dynamic game when each firm's strategy specifies a rate of production at any moment as a function of the joint capacities for future production remaining as of that moment. An equilibrium is a pair of such state–contingent rules which are optimal for each firm, given the other's strategy, from every conceivable state. In a sense our game is a "closed–loop", "multiple stock" version of the "oil"igopoly problem encountered in the exhaustible resources literature.[1] This is, to our knowledge, the first treatment of perfect equilibrium in a depletable resource model with multiple stocks. Previous work has succeeded in imposing subgame perfection only in the context of "common pool" situations, in which there is a single limiting resource.[2]

The assumed indivisibility in production further distinguishes the present analysis from existing work in resource economics. Our firms decide whether or not to produce at each moment, but are unable to vary the non-zero rate at which output is supplied. They regulate an "on–off" control variable, facing a constraint on the total amount of time this variable can be "on" over the horizon of the game. In effect, their output decisions are

---

[1] See Loury (1986), and the literature cited there.

[2] An early contribution was the study of competition in the exploitation of a common fishery by Levhari and Mirman (1980). Related analyzes include Reinganum and Stokey (1985), and the recent, more general, work of Sundaram (1989) and Dutta and Sundaram (1989).

qualitative choices about participation. The game is analogous to a dynamic Prisoner's Dilemma in which the capacity to play "tough" in the future diminishes with an increased frequency of having played "tough" in the past.[3] We focus on the classic question in the study of the Prisoner's Dilemma — the extent to which, in a non–cooperation equilibrium, the players can avoid dissipative competition through mutual forebearance.

Characterization of closed–loop equilibrium in non–zero sum differential games with several state variables is a very formidable, largely open, problem. Explicit solutions are rare outside of the linear–quadratic context, though a number of economically interesting examples have been identified with a special structure permitting profitable study.[4] In this paper we make progress by simplifying the problem in two ways: As noted, we limit the set of actions from which the players can choose at any instant; and, we assume that actions, once chosen, must remain fixed for some (possibly quite short) interval of time. These assumptions permit explicit solutions for perfect equilibria to be obtained.

Our ability to solve the model explicitly is exploited to pose a number of interesting questions for which general answers are unavailable in the existing literature on dynamic duopoly: Are firms more, or less, able to avoid mutually destructive competition when they can, and do, commit at the initial date to specific future actions, relative to a situation in which such commitment is not possible? What is the relationship in equilibrium between the size and distribution of payoffs, on the one hand, and capacities to produce, on the other? How will the distribution of remaining capacity evolve over time? How are

---

[3]A closely related model is the continuous time Prisoner's Dilemma game of Rosenthal and Spady (1989). There Markov Perfect Equilibria in strategies defined on a two dimensional state space of cumulative profits are studied.

[4]A good discussion of the differential game approach to dynamic oligopoly problems is in Fudenberg and Tirole (1986). Two linear–quadratic analyses are Fershtman and Kamien (1987), who study a deterministic, one-state-variable model of duopolistic pricing, and Hansen, Epple and Roberds (1985), who solve a stochastic two–state-variable model of resource depletion, though without imposing sugbgame perfection. Other models with a structure permitting profitable study of subgame perfect equilibrium are the memoryless patent race of Reinganum (1981), the racing model of Harris and Vickers (1985), the capacity investment model of Spence (1979), extended in Fudenberg and Tirole (1983), and the duopoly pricing model with menu costs of Halperin (1990).

these considerations affected by the rate of discount and the elasticity of market demand?

Among the results obtained are: (1) state—contingent strategies imply equilibrium payoffs at least as great, and often greater, than those possible with open—loop strategies; (2) symmetry of capacities facilitates collusion in closed—loop equilibrium, but can undermine it when open-loop strategies are employed; (3) the extent of endowments consistent with attaining tacit collusion in perfect equilibrium is inversely related to total capacity, to the interest rate, and to the elasticity of demand; (4) the ability to credibly threaten future production is a basic strategic advantage which allows the appropriation of a more than proportionate share of the industy's total payoff. Result (1) implies that the ability to commit at the initial date to future actions can be disadvantageous! Result (2) suggests that, when such commitment is not possible, firms are better able to maintain high prices to their mutual advantage if neither of them enjoys a large market share! Result (4) provides theoretical support for the common intuition that a firm with a substantial share of total capacity (Saudi Arabia in OPEC!) can disproportionately influence the course of play.

To obtain these results we employ the two simplifications mentioned to transform a complete information, continuous time game of dynamic duopoly into a more tractable discrete time analogue. This conversion is described in section II. Perfect equilibrium in Markov strategies is defined in section III. A two–dimensional backward induction is employed in section IV to solve for explicit equilibria, which are compared to the open-loop Nash equilibria of the same game. In section V we consider the limit of equilibrium behavior in the discrete model as period length shrinks, raising some technical and conceptual issues discussed briefly in that and the concluding section of the paper.[5]

---

[5]Because of the indivisibility in production, coordination of play is closely linked with the idea of "precedence" in this model. Our discrete time equilibria exploit this, and the related notion of "alternation" heavily. In continuous time, when players are free to change their actions at any instant, it is nolonger clear how to define who produced "last period". Thus, it is not clear how to interpret the limit, as period length goes to zero, of the discrete time equilibria analyzed in section IV. See the treatment of extensive form games in continuous time by Simon and Stinchcombe (1989) for a fuller discussion of these issues.

II. The Model

We consider the following model of competition between two independent firms, A and B, facing an indivisibility in production, with given limitations on their cumulative capacities to produce. Time is continuous over an infinite horizon, $\tau \in [0, \infty)$. At date τ the flow rates of production of firms A and B, denoted by q_τ^a and q_τ^b respectively, are such that $q_\tau \equiv (q_\tau^a, q_\tau^b) \in \{0,1\}^2$; that is, either a firm is inactive, or it is producing at a unitary rate at any moment. $Q_\tau = q_\tau^a + q_\tau^b$ is the industry's production rate at τ; obviously $Q_\tau \in \{0,1,2\}$, $\forall \tau \geq 0$. A firm's *initial capacity* (or *stock*) is a number R_j (j=A,B) denoting feasible cumulative (resource cost of) production for that firm over the course of the game; $R \equiv (R_A, R_B) \in \mathbb{R}_+^2$ is the vector of initial capacities/stocks. Capacity falls at a unitary rate if a firm is active; a firm exits the game when its capacity has fallen to zero. There is no way to add to stocks during the play of the game. The firms' stocks are common knowledge.

The demand side of the market is passively modelled; buyers do not behave strategically. There is an inverse demand function, $P(\cdot)$, which is time invariant and dependent only on the total rate of flow of output of the two firms. Monetary units are such that $P(1)=1 > \epsilon \equiv P(2) > 0$. (For completeness, let $P(0)$ be any finite number.) If $\epsilon < 1/2$ then industry flow revenue falls as the rate of production increases, so we say that demand is *inelastic*. If $\epsilon \geq 1/2$, then we will refer to demand as being *elastic*. There is no other cost of production beyond the opportunity cost implied by the depletion of capacity. Future receipts are discounted at the common, uniform rate $r>0$. The firms' payoffs are simply their respective discounted flows of revenue received over the infinite horizon as a result of their joint paths of supply.

We study the equilibria of the following discrete time analogue of the continuous time set-up introduced above: Assume that the firms cannot change their flow rate of output except at particular, exogenously given and uniformly spaced moments in time — that is, except at the beginning of a "period." After any change they are bound to maintain their production at the new rate until the next such moment. These moments when

adjustments are permitted are common knowledge to the two firms. Let h>0 be the length of a "period" — the interval between dates when a change of action is allowed. A discrete time dynamic game is defined for each value of h, whose equilibria we characterize. We also study the behavior which emerges in these equilibria in the limit, as h→0.

Suppose then that $T(h) \equiv \{0, h, 2h, 3h,...\}$ is the set of dates when new actions can be taken. Associate with each $\tau \in T(h)$ the integer $t = \tau/h$. We will often refer to "date t", meaning the instant $\tau = h \cdot t$, and to "period t", meaning the interval $[h \cdot t, h \cdot (t+1))$. Define the *discount factor* $\delta(h) \equiv e^{-rh}$; a dollar received on date t is worth $\delta(h)^t$ dollars at date zero. Monetary flows over intervals $[\tau, \tau+h)$, $\tau \in T(h)$, are constant given our assumptions, and may be regarded as equivalent lump sums paid at the beginning of each period. Suppose that at dates $t \in I$ the firms take the actions q_t. Then their respective payoffs are V_A and V_B where:

(2.1) $$V_A = \beta \sum_{t \in I} \delta^t [q_t^a \cdot P(Q_t)]; \quad \text{and} \quad V_B = \beta \sum_{t \in I} \delta^t [q_t^b \cdot P(Q_t)],$$

for $\beta(h) \equiv [\frac{1-e^{-rh}}{r}]$, the lump sum equivalent of the flow of one dollar over an interval of length h. Unique payoffs are implied by every path of firms' actions via (2.1).

Given periods of length h, we consider now how the firms' productive capacities evolve over the course of the game. Let $\bar{I}$ be the set of positive integers, and $I \equiv \bar{I} \cup \{0\}$. Define the *initial state*, $s_0(h,R) \equiv (a_0(h,R_A), b_0(h,R_B))$, as follows: $a_0 \equiv \max\{a \in I \mid a \cdot h \le R_A\}$, and $b_0 \equiv \max\{b \in I \mid b \cdot h \le R_B\}$. (We suppress explicit dependence of s_0 on h and R when that is inessential.) The initial state is a pair of integers denoting the "number of plays" which each firm can make over the course of the game, where a "play" involves supplying output at the constant unitary rate for a period of duration h, and where we "round–off" by neglecting the fact that the initial stocks may not be evenly divisible by h. Let $S \equiv I \times I$ denote the *state space*. In this way, for every initial vector of capacities $R \in \mathbb{R}_+^2$ and every

$h > 0$, we assign an initial state $s_0 \epsilon S$, representing this association by: $s_0(h,R) \approx h^{-1} \cdot R$. It is apparent that if s_t is the *state of the game at date* t, and if q_t is the vector of firms' actions during period t, then the state at date t+1 satisfies: $s_{t+1} = s_t - q_t$.

III. Markov Perfect Equilibrium

It is obvious that, for any given history of play prior to date t, the set of feasible action paths subsequent to date t, and the firms' conditional valuations of these paths, depend on that history only through the state of the game at t, $s_t \equiv (a_t, b_t)$. That is, the "state of the game", as defined here, is a "payoff relevant state" in the sense, e.g., of Maskin and Tirole (1989): Conditional valuations are independent of the history, given the stationarity assumed in (2.1). Feasibility depends only on whether remaining stocks are big enough; all histories prior to t which reach the state s_t imply the same set of feasible subsequent actions.

We employ the concept of Markov Perfect Equilibrium (MPE) for the discrete game, with h fixed, in the analysis to follow. In general then, strategies are functions from S into [0,1], reporting for each firm the probability of participation in the market during a period when the state is $s \epsilon S$. However, we shall restrict attention to pure strategy MPE. This restriction involves less loss of generality than may appear to be the case, since pure strategy equilibria always exist in this model, and mixed strategies are inefficient here. If, at some node s, A and B play "0" with positive probability, they risk getting "stuck" there for a period — a source of lost industry profits which does not exist with pure strategies. Because we are primarily concerned with understanding when a revenue maximizing outcome occurs in equilibrium, we restrict attention to pure strategies.

Given a pair of pure Markov strategies on S, the resulting motion through the lattice during the play of the game may be envisioned as exemplified by Figure 1. At each point $s \epsilon S$ four continuations are possible, corresponding to the four elements $q \epsilon \{0,1\}^2$. These outcomes result in motion of the state from s to some successor node, $s' = s - q$. If,

at node s, both firms stay out of the market (i.e., q=(0,0)), then s′=s, and the state remains where it began. With Markov strategies this amounts to remaining at s forever, an obvious impossibility in an equilibrium unless s=(0,0). The three other continuations correspond to one or both firms being in the market, and result in motion of the state either "due South", "due West", or "Southwest", as depicted by the arrows in the Figure. Eventually the state reaches one or both of the axes, representing the exit of one or both firms. At this point, the firm with remaining capacity produces one unit in each period until the state reaches the origin.

Figure 1

Thus, with each profile of pure Markov strategies and each initial state may be associated a unique path through the lattice and, via (2.1), unique payoffs for the firms. Given the period length, these payoffs are a function of the initial node, the strategy pair, and the parameters ϵ and δ. Every path terminates in a finite number of steps. Current period payoffs (divided by $\beta(h)$) at each node depend on the actions taken there as follows:

Table 1

We now procede formally, introducing some notation and definitions. Define:

$$\Sigma^A \equiv \{\sigma_A : S \to \{0,1\} \mid \sigma_A(0,b)=0, \forall b \in I\}, \quad \Sigma^B \equiv \{\sigma_B : S \to \{0,1\} \mid \sigma_B(a,0)=0, \forall a \in I\}, \text{ and } \Sigma \equiv \Sigma^A \times \Sigma^B.$$

A *pure strategy profile* (or simply, a profile) is an element $\sigma \epsilon \Sigma$. From any state $s \epsilon S$, every profile $\sigma \epsilon \Sigma$ generates a path $\{s_0, s_1, ..., s_t, ...\}$, defined inductively as follows:

(3.1)
$$s_0 = s, \quad \text{and}$$
$$s_t = s_{t-1} - \sigma(s_{t-1}), \quad t \epsilon I.$$

We denote this sequence $\{s_t(s, \sigma)\}$. We say that point s' *is a successor of* s *under* σ if $s' \epsilon \{s_t(s, \sigma)\}$, and that s' *is a predecessor of* s *under* σ if $s \epsilon \{s_t(s', \sigma)\}$. Identifying q_t in (2.1) with $\sigma(s_t(s, \sigma))$, we may define the *value function* for each firm, $V_A(\sigma)(s)$ and $V_B(\sigma)(s)$, at state $s \epsilon S$ given profile $\sigma \epsilon \Sigma$, in the obvious way, as the respective sum in (2.1) along the path $\{s_t(s, \sigma)\}$.

Definition 3.1: A strategy profile $\hat{\sigma} \epsilon \Sigma$ is a *Nash Equilibrium from* $s \epsilon S$ if:
$$V_A(\hat{\sigma})(s) \geq V_A(\sigma_A, \hat{\sigma}_B)(s) \quad \forall \sigma_A \epsilon \Sigma^A, \text{ and}$$
$$V_B(\hat{\sigma})(s) \geq V_B(\hat{\sigma}_A, \sigma_B)(s) \quad \forall \sigma_B \epsilon \Sigma^B.$$

Definition 3.2: A strategy profile $\sigma^* \epsilon \Sigma$ is a *Markov Perfect Equilibrium* if it is a Nash Equilibrium from s, for every $s \epsilon S$.

Let $\Sigma^* \subset \Sigma$ be the set of all MPE's in pure strategies. Nash Equilibrium is weaker than MPE since it requires strategies to be best responses to each other only from a particular initial state. In a NE neither firm gains by deviating from its strategy at any node along the subsequent path, given that it anticipates the response to the deviation to be as reported by the strategy of the other firm. Yet these responses are not themselves constrained, beyond feasibility, by the definition. Therefore, NE may depend on threats which are not credible. MPE requires that promised responses to deviations must be consistent with subsequent equilibrium play.

We now formalize the idea that a profile σ is an MPE on some subset of the state space. In the sequel we construct equilibria by finding particular strategies that are best responses to each other from initial states in a given region of S, extending the resulting profile to the entire state space so as to constitute a full equilibrium. We say that a set of states W is *closed to deviations* from a profile σ if all paths from $s \in W$ generated by σ stay in W, and if one firm's deviation from such a path also yields a point in W. Denote by $\sigma \backslash \bar{\sigma}_j$ the profile obtained by replacing j's strategy under σ with $\bar{\sigma}_j$, $j = A, B$.

Definition 3.3: Given a set of states $W \subset S$ and a function $\sigma: W \to \{0,1\}^2$, we say that W is *closed to deviations from* σ if the following condition holds:

$$\forall s \in W, \forall \bar{\sigma}_A \in \Sigma^A, \forall \bar{\sigma}_B \in \Sigma^B: \{s_t(s, \sigma \backslash \bar{\sigma}_A)\} \subset W, \text{ and } \{s_t(s, \sigma \backslash \bar{\sigma}_B)\} \subset W.$$

Definition 3.4: A strategy profile $\sigma \in \Sigma$ is a *restricted MPE on* W if:

    (a) W is closed to deviations from σ; and

    (b) σ is a NE from s, $\forall s \in W$.

Let $\Sigma^*_{|W} \subset \Sigma$ denote the class of restricted MPE on W. Obviously, if $\sigma \in \Sigma^*_{|W}$, then $\sigma' \in \Sigma$ and $\sigma' \equiv \sigma$ on W implies $\sigma' \in \Sigma^*_{|W}$. Restricted equilibria on a given set may thus be thought of as equivalence classes of strategy profiles, in which membership is defined by agreement on the set in question. It is also obvious that if $\sigma \in \Sigma^*$ then $\sigma \in \Sigma^*_{|W}$ if and only if W is closed to deviations from σ. Theorem 1 establishes that every restricted MPE can be extended to the entire state space in a manner consistent with Markov Perfection.

Consider the following simple example which illustrates the definitions. For $n \in I$ define $S^n \equiv \{s = (a,b) \in S \mid a+b \leq n\}$. Obviously, S^n is closed to deviations from any $\sigma \in \Sigma$ (as is any set containing the origin, and containing all points "southwest" of any point it contains.) Now, consider the strategy profile $\bar{\sigma} = (\bar{\sigma}_A, \bar{\sigma}_B)$ defined as follows: B plays whenever he has a positive stock. A waits until B has exhausted before playing any-

thing. Let $n' \equiv \text{Max}\{n \in I \mid \delta^n \geq \epsilon\}$. Then, if $n' \geq 1$, we have that $\overline{\sigma} \in \Sigma^*_{\mid S^n}$ for $2 \leq n \leq n'+1$: B cannot gain from deviating; and A will deviate only if $\epsilon > \delta^n$. But it is easy to see that $\overline{\sigma} \notin \Sigma^*$. For $\max\{a,b\}$ large enough, A will deviate.

To establish that a profile σ constitutes an MPE it suffices to show that, for every $s \in S$, neither firm can increase his payoff by a single deviation from σ at some point along the path $\{s_t(s,\sigma)\}$, given that σ is followed by both firms after the deviation. That is, only "one–shot" deviations need be considered. This follows from the obvious fact that, if a sequence of deviations is desirable, it must contain a "one–shot" deviation at some node which raises the firm's payoff. Moreover, since no firm will deviate at a node where he alone produces, only two forms of deviation are relevant: Either $\sigma(s)=(1,1)$ and someone prefers to stay out of the market, or $\sigma(s) \in \{(1,0),(0,1)\}$ and the firm not producing prefers to be active.

At the node s, for the firm A, and given the putative equilibrium profile σ, to determine if a profitable deviation is possible we need to compare two quantities: (i) $\beta\epsilon + \delta V_A(\sigma)(s')$, the current revenue flow when $q=(1,1)$ plus the discounted continuation value of moving from s to $s' \equiv s-(1,1)$; and (ii) $\delta V_A(\sigma)(s'')$, the discounted continuation value occasioned by the actions $q=(0,1)$ which lead to the node $s'' \equiv s-(0,1)$. If $\sigma(s)=(1,0)$, A never deviates from σ at s. If quantity (i) exceeds (ii) and $\sigma(s)=(0,1)$, or if (ii) exceeds (i) and $\sigma(s)=(1,1)$, then A deviates. The analogous conditions characterize the desirability of deviations for firm B.

If a profile σ provides no incentive for either firm to deviate from its specified actions at a given node, assuming it is followed thereafter, then the firms' value functions $V_A(\sigma)(s)$ and $V_B(\sigma)(s)$ must satisfy standard dynamic programming conditions. Thus:

Definition 3.5: Given strictly positive $s \in S$, we say that a strategy profile $\sigma \in \Sigma$ *satisfies the Bellman Inequalities* at s if $\sigma(s) \neq (0,0)$, and:

(3.2) (a) $\sigma(s)=(1,1) \Rightarrow V_A(\sigma)(s) = \beta\epsilon + \delta V_A(\sigma)(s-(1,1)) \geq \delta V_A(\sigma)(s-(0,1)),$

and $V_B(\sigma)(s) = \beta\epsilon + \delta V_B(\sigma)(s-(1,1)) \geq \delta V_B(\sigma)(s-(1,0)).$

(b) $\sigma(s)=(0,1) \Rightarrow \beta\epsilon + \delta V_A(\sigma)(s-(1,1)) \leq \delta V_A(\sigma)(s-(0,1)) = V_A(\sigma)(s).$

(c) $\sigma(s)=(1,0) \Rightarrow \beta\epsilon + \delta V_B(\sigma)(s-(1,1)) \leq \delta V_B(\sigma)(s-(1,0)) = V_B(\sigma)(s).$

We summarize the foregoing definitions and discussion with the following useful result:

Lemma 1: Given $W \subset S$ let the profile $\sigma \in \Sigma$ satisfy the following conditions on W:

(a) W is closed to deviations from σ;

(b) $\sigma_A(a,0)=\sigma_B(0,b)=1$, at any non-zero points $(a,0)$ and $(0,b) \in W$;

(c) σ satisfies the Bellman Inequalities at every strictly positive $s \in W$;

Then $\sigma \in \Sigma^*_{|W}$. If σ satisfies (b) and (c) on all of S, then $\sigma \in \Sigma^*$.

It is possible, using a two–dimensional backward induction, to characterize all Markov Perfect Equilibria in pure strategies for this game. This may be seen as follows: For arbitrary values of V_A and V_B at the nodes $s-(1,1)$, $s-(1,0)$ and $s-(0,1)$, there is always some pair of actions $q=\sigma(s)$ which permit the Bellman Inequalities to be satisfied at s. That is, for s strictly positive and $\sigma \in \Sigma$, either both inequalities in (3.2)(a) hold or one of them fails, in which case one of the inequalities in (3.2)(b) or (3.2)(c) holds. Moreover, $V_A(\sigma)(n,0)=V_B(\sigma)(0,n)=\beta(1-\delta^n)/(1-\delta)$ for all σ satisfying (b) of Lemma 1. Therefore, starting at the point $(1,1)$ and working systematically "backward", up and to the right in the lattice, one constructs a pure strategy MPE by successively assigning values to σ at s, given the values already assigned at to σ at $s' < s$, in such a way that the Bellman Inequalities are satisfied at each step. It is possible to do this so that at each node s, values for the functions $V_A(\sigma)(\cdot)$ and $V_B(\sigma)(\cdot)$ have already been determined at the nodes $s-(1,1)$, $s-(1,0)$ and $s-(0,1)$. Employing this induction, we prove that a restricted MPE on a subset W of the state space can always be extended to the entire lattice in such

a way that the resulting profile constitutes a full equilibrium on S.

Theorem 1: Given $W \subset S$ and $\bar{\sigma} \in \Sigma^*_{|W}$, there exists $\sigma^* \in \Sigma^*$ such that: $\sigma^*(s) = \bar{\sigma}(s)$, $\forall s \in W$.
Proof: See the apppendix.

Theorem 1 is an important tool which shall be used repeatedly in the sequel. We turn now to a substantive analysis of MFE strategy profiles. The argument is constructive; we solve the model explicitly to find state–contingent plans of action which are mutually consistent with revenue maximization by both firms from every conceivable state. We begin by defining tacit collusion and describing open–loop equilibrium in the model.

IV. On the Possibility of Collusion in Markov Perfect Equilibria

A. Tacit Collusion

In this simple setting, with only two possible actions, the notion of tacit collusion takes a very primitive form. In order to maximize the sum of their payoffs the firms must arrange to avoid supplying at the same time if that reduces discounted industry revenues. It is obvious that the industry revenue maximizing output at any strictly positive node $s = (a,b)$ depends only on the total $(a+b)$. Let $Q^*(s) \in \{0,1,2\}$ be this optimal collusive behavior at $s \in S$. Regarding tacit collusion (or coordination, or cooperation), we say that the MPE $\sigma \in \Sigma^*$ *achieves collusion at* s if $\sigma_A(s) + \sigma_B(s) = Q^*(s)$. The *region of collusion* under σ, $M(\sigma)$, is the set of nodes at which σ achieves collusion. If a path satisfies $\{s_t(s,\sigma)\} \subset M(\sigma)$, $s \in S$ and $\sigma \in \Sigma^*$, we call it a *collusive path*. Industry revenues are maximized in the MPE σ from the initial node s_0 iff σ generates a collusive path from s_0.

With inelastic demand it is obvious that $Q^*(s) \equiv 1$, $s \neq (0,0)$, since supplying two units instead of one lowers both current and future revenue. With elastic demand, if $s = (a,b)$ and $a+b$ is "large", then $Q^*(s) = 2$: Because of discounting, the opportunity cost of supplying another unit today vanishes as total capacity grows large, while the marginal rev-

enue of two instead of one unit of output is $\beta(2\epsilon-1)>0$. So if $\epsilon>1/2$ then, for some $\bar{n}\geq1$, $Q^*(a,b)=1$, $a+b\leq\bar{n}$; and $Q^*(a,b)=2$, $a+b>\bar{n}$. The critical $\bar{n}$ is given as follows:

$$(4.1) \qquad \bar{n} = \text{Max}\{n\epsilon I| \frac{1-\delta^n}{1-\delta} \geq 2\epsilon+\epsilon(\frac{1-\delta^{n-2}}{1-\delta})\} = \text{Max}\{n\epsilon I| \delta^{n-1} \geq 2\epsilon-1\}.$$

That is, $\bar{n}$ is the largest total stock at which the opportunity cost of an incremental unit of production, given that one unit is to be supplied in each subsequent period, is no less than the marginal revenue of raising output from one to two units. With inelastic demand, collusion in equilibrium requires the avoidance of simultaneous production everywhere. With elastic demand, such coordination is only relevant when the total stock is less than $\bar{n}$.

B. Collusion in Open Loop Equilibrium

In order to contrast the possibilities for collusion in Markov Perfect Equilibria with the outcomes achieved in the open-loop (precommitment) equilibria of the same game, we consider the latter briefly. By "open—loop" play we mean that firms can bind themselves at the initial date to time paths of actions which are known to both. If one firm believes the other is so committed, then it has a corresponding path which maximizes its payoff. With paths constrained to be compatible with the available capacities of the firms, we consider the Nash Equilibria in these paths. Such an equilibrium achieves tacit collusion from the node (a,b) if it avoids simultaneous production when to do so is necessary to maximize the sum of the firms' payoffs. The following result characterizes the region in S where tacit collusion may be achieved using open-loop strategies.

Lemma 2: There exists an open—loop Nash Equilibrium from the strictly positive node (a,b)ϵS which avoids simultaneous production everywhere along its path if and only if:

(4.2)
$$\tfrac{1}{2}\text{Max}\{a,b\} + \text{Min}\{a,b\} \leq n' \equiv \text{Max}\{n\epsilon I \,|\, \delta^n \geq \epsilon\}.$$

Proof: In an open—loop Nash Equilibrium neither firm wishes to shifts a unit of output from its last date of positive supply to its first date of zero supply. Letting the difference in these two dates for either firm be θ, such a shift pays if $\epsilon > \delta^\theta$. So an open—loop equilibrium without simultaneous production is impossible if $\epsilon > \delta^\theta$ for some firm, along every path where one or the other, but not both of them, supplies on successive dates. Consider the problem of finding a sequence of a "1's" and b "0's" such that the larger of: (1) the difference in order between the last "1" and the first "0"; and (2) the difference in order between the last "0" and the first "1"; is minimal. Denote this minimal difference by $\underline{\theta}(a,b)$. Thinking of this sequence as the time path of supply for firm A in an open-loop equilibrium from initial node (a,b) which avoids simultaneous production, we see that neither firm will deviate if $\underline{\theta}(a,b) \leq n'$. The reader can verify that the minimal difference $\underline{\theta}(a,b)$ is realized by placing the lesser in number of the two digits in the middle of the sequence and surrounding them as evenly as possible on either side by the other digit. Assuming that $a \geq b$, and that a is even, this is illustrated in Figure 2:

$$\underbrace{1,\ 1,\ 1,\ldots 1}_{a/2},\ \underbrace{0,\ 0,\ 0,\ldots 0}_{b},\ \underbrace{1,\ 1,\ 1,\ldots 1}_{a/2}$$

Figure 2

Examination of this Figure shows that this disposition cannot be improved upon, and that
$$\underline{\theta}(a,b) = \text{Min}\{n\epsilon I \,|\, n \geq \tfrac{1}{2}\text{Max}(a,b)+\text{Min}(a,b)\}. \qquad \square$$

With open—loop strategies each firm takes the other's path as given, and so anticipates no consequence of an increase in current supply besides a reduction of its own capacity for future production. So the "shadow price" on a marginal unit of capacity is the present value of revenue from the last unit sold. There exist endowments from which each

firms' initial shadow price on marginal capacity is less than ϵ, its marginal revenue at zero output, along every collusive path. From such endowments, for which (4.2) fails, both firms produce on the initial date in every open—loop equilibrium. (Note: $2n' \leq \bar{n}$.) So, whatever the elasticity of demand, total size and relative symmetry of capacities can militate against achieving maximal industry revenues in an open—loop equilibrium.

C. Precedence and Alternation with Markov Strategies

We will show that something close to the opposite of the foregoing is true regarding collusion in Markov Perfect Equilibrium. Our method is constructive; we shall exhibit a particular equilibrium profile to establish this point. Intuitively, closed-loop strategies are more effective at achieving coordination because, by making the entire future course of the game contingent upon current actions, they can impose on the firms a cost of increased output far greater than the foregone discounted revenue from the last unit sold. We demonstrate formally in Proposition 1 that when total production capacity is arbitrarily large, implying a vanishingly small shadow price for the marginal unit, if capacities are relatively symmetric and the periods short, the firms can always coordinate their actions tacitly so as to maximize discounted industry revenue in a non—cooperative, perfect equilibrium.

A key concept in this construction is the notion of "turn—taking". With $\{0,1\}$ actions and Markov strategies, colluding to maintain a high price amounts to specifying regions of the state space within which one firm has precedence, and the other is expected to forebear. Along a collusive path, one firm produces until a node is reached which signals the other's "turn", the other then produces until the state passes out of its region of precedence, and so on. In equilibrium the firm without the turn must prefer forebearance to "playing out of turn."

A special case of this turn-taking partition might be called "pure alternation," where the firms share the market by producing on alternate dates. To motivate the alternating strategy profile, consider a one—shot game with payoffs as given in Table 1.

This game has (1,1) as its strictly dominant Nash Equilibrium, providing each firm with a return of ϵ. Suppose that the firms were to encounter this game over an infinite sequence of periods, with δ being their common discount factor. For small ϵ the one-shot equilibrium would imply substantial losses. Both firms would be much better-off if they could coordinate their actions so as to avoid "head to head" competition, yet each prefers to supply in the current period. Resolution of this Prisoner's Dilemma—like problem requires that the firms reach some tacit agreement as to who will supply, and when.

A natural understanding which might develop between them in such a situation is: "Whoever produced last period stays out this period". If the discount factor is not too low it will be better to wait until tomorrow for one unit of revenue, and the implied opportunity to wait yet again the day after tomorrow, than to claim ϵ units of revenue every day. It is easy to see that such an *alternating convention*, supported by the threatened punishment of reverting to "(1,1) forever" if deviation occurs, is a subgame perfect equilibrium of the infinitely repeated game if and only if "waiting to alternate" is better than "competing forever". That is, if and only if: $\delta \cdot (1 + \delta^2 + \delta^4 + ...) = \delta/(1-\delta^2) \geq \epsilon/(1-\delta)$. Or,

$$(4.3) \qquad\qquad \epsilon \leq \frac{\delta}{1+\delta}.$$

This condition is important in the analysis of Markov Equilibrium to follow. Note that it holds with inelastic demand if the periods are short enough, but never with elastic demand.

Return now to the dynamic game with cumulative capacity constraints which we have been considering. Formally, define a profile $\sigma \in \Sigma$ as *an alternating strategy profile* if, for all strictly positive $s \in S$: $\sigma(s) + \sigma(s - \sigma(s)) = (1,1)$. The path $\{s_t(s,\sigma)\}$ therefore involves the firms supplying the market alternately in successive periods for as long as both have positive stocks. There are exactly two such profiles, distinguished by which of the firms supplies when the stocks of both are positive and their sum is an even integer. We assume without loss of generality that firm A produces under the alternating strategy profile

when the sum of positive stocks is even, and denote the corresponding profile by $\hat{\sigma} \epsilon \Sigma$. Notice how the reduction of the continuous game to one in discrete time with a countable state space, induced by the assumption that actions cannot be instantly changed, is being heavily exploited in defining this form of precedence.

Figure 3

Take an arbitrary point $s \epsilon S$, and consider the set of all successor nodes to s under $\hat{\sigma}$, which we call the *alternating path from* s, denoted $\{s_t(s,\hat{\sigma})\}$. These paths have a stair-step appearance, and include nodes along the axes leading to $(0,0)$. Figure 3 illustrates several alternating paths. Notice that these paths, regarded as subsets of S, are closed to deviations from $\hat{\sigma}$, $\forall s \epsilon S$. A deviation from $\hat{\sigma}$ means either remaining at s, or "jumping ahead" to $s-(1,1)$. By definition of $\hat{\sigma}$, $s-(1,1)=[s-\hat{\sigma}(s)-\hat{\sigma}(s-\hat{\sigma}(s))]=s_2(s,\hat{\sigma})\epsilon\{s_t(s,\hat{\sigma})\}$, for $s>>(0,0)$, so a deviation from $\hat{\sigma}$ always leads to a node belonging to its alternating path through s. (Indeed, <u>only</u> an alternating strategy profile generates paths which are closed to deviation from itself.) A useful way to state this fact is that, along any alternating path, the consequence of deviation from "waiting one's turn" is always to "lose one's turn."

It is easy to see that the turn-taking convention represented by $\hat{\sigma}$ is not an MPE. For if $\delta^a < \epsilon$ it must be that $\sigma_A(a,1)=1$, $\forall \sigma \epsilon \Sigma^*$: By playing "1" at $(a,1)$ A guarantees itself a payoff of at least $\beta[\epsilon+\delta(1-\delta^{a-1})/(1-\delta)]$ which exceeds $\beta\delta(1-\delta^a)/(1-\delta)$, the best it can get from playing "0." Since $\hat{\sigma}_A(a,1)=0$ when a is even, it follows that $\hat{\sigma} \notin \Sigma^*$.

Nevertheless, there is some non-trivial subset of nodes, implicitly defined as follows, from which the alternating profile $\hat{\sigma}$ is a <u>restricted</u> MPE:

Definition 4.1: For $\hat{\sigma}$ as above and $0 < \epsilon < 1$ define the *region of alternation*, $\mathscr{D}(\epsilon)$, by:

$$\mathscr{D}(\epsilon) \equiv \{s \in S \mid \forall s' \in \{s_t(s,\hat{\sigma})\}, \text{ if } \hat{\sigma}_j(s')=0, \text{ then } V_j(\hat{\sigma})(s') \leq \delta \cdot \beta[\tfrac{1-\epsilon}{1-\delta}], \, j=A,B\}.$$

Theorem 2: $\hat{\sigma} \in \Sigma^*_{|\mathscr{D}(\epsilon)}$. Moreover, if $\hat{\sigma} \in \Sigma^*_{|W}$ then $W \subset \mathscr{D}(\epsilon)$. That is, $\mathscr{D}(\epsilon)$ is the largest subset of S on which a mutual agreement to "take turns" could be self–enforcing.

Proof: $\mathscr{D}(\epsilon)$ is the countable union of alternating paths, each of which is closed to deviations from $\hat{\sigma}$. Therefore $\mathscr{D}(\epsilon)$ is closed to deviations from $\hat{\sigma}$. The Theorem then follows from Lemma 1, if we can show that Bellman Inequalities (3.2)(b or c) are satisfied at every strictly positive node $s \in \mathscr{D}(\epsilon)$. If firm j "waits its turn" at s its payoff there is $V_j(\hat{\sigma})(s)$. If it deviates by playing "out of turn" it receives $\beta\epsilon$ immediately, but loses the opportunity to have the market to itself next period. This continuation yields the payoff $\beta\epsilon + \delta V_j(\hat{\sigma})(s-(1,1))$. But, because it is waiting to have the market to itself, it must be that $V_j(\hat{\sigma})(s)=\delta \cdot [\beta + \delta V_j(\hat{\sigma})(s-(1,1))]$. So the deviation pays iff $V_j(\hat{\sigma})(s) > \delta\beta(\tfrac{1-\epsilon}{1-\delta})$. Thus, $\mathscr{D}(\epsilon)$ is, by definition, precisely the set of nodes where deviation from $\hat{\sigma}$ does not pay. $\square$

This result leads to the following general rule about turn–taking: <u>A firm "waits its turn" for one period if the value of that turn, when it does come, does not exceed the value of receiving, in perpetuity, the revenue loss to the other firm caused by it having played out of turn.</u> Moreover, the argument for this general rule only used the alternating property of $\hat{\sigma}$ "locally," at a node s and its immediate successor. So we have the following partial converse of Theorem 2, relevant whenever precedence changes hands.

Corollary 1: Suppose $\sigma \epsilon \Sigma^*$. Then $\sigma(s) + \sigma(s-\sigma(s)) = (1,1)$, and $\sigma_j(s)=0$ imply that $V_j(\sigma)(s) \leq \delta\beta(\frac{1-\epsilon}{1-\delta})$, for $j \epsilon \{A,B\}$ and $s \epsilon S$.

The next result provides an explicit description of $\mathscr{D}(\epsilon)$. We also introduce notation which will be used extensively in the sequel. Recall that $n' \equiv \text{Max}\{n \epsilon I \mid \delta^n \geq \epsilon\}$.

Definition 4.2: For $s=(a,b) \epsilon S$ define $n(s) \equiv |a-b|$, and $k(s) \equiv \min\{a,b\}$. Let $j(s)$ denote the larger firm and $i(s)$ the smaller firm at s, $i,j \epsilon \{A,B\}$. (Explicit dependence of these functions on s will often be suppressed.) Define the function $\phi(n,k)$ as follows:

$$(4.4) \qquad \phi(n,k) \equiv \delta^{2k-2} \cdot \delta^{\bar{n}} + (1-\delta^{2k-2}) \cdot (\frac{\delta}{1+\delta}), \text{ for } n \epsilon I \text{ and } k \epsilon I.$$

Corollary 2: $\mathscr{D}(\epsilon)=\{s \epsilon S \mid \forall k \epsilon \bar{I}, k \leq k(s): \phi(n(s)+1,k) \geq \epsilon$ if $\hat{\sigma}_{j(s)}(s)=0$, and $\phi(n(s),k) \geq \epsilon$ if $\hat{\sigma}_{i(s)}(s)=0\}$. Moreover, when (4.3) holds then $\mathscr{D}(\epsilon)=\{s \epsilon S \mid \forall t \geq 0, n(s_t(s,\hat{\sigma})) \leq n'\}$.
Proof: See the appendix.

With indivisible production forebearance, by one firm or the other, is essential if they are to share the market and maintain a high price. Theorem 2 and Corollary 1 imply that such restraint is consistent with the firms' interests only if the "contingent right to produce" has been allocated in such a way that payoffs never exceed $\beta(\frac{1-\delta}{1-\epsilon})$. Any greater return is too large for a firm to be willing to "wait its turn" to receive it. Now the region $\mathscr{D}(\epsilon)$ consists of nodes from which this criterion is satisfied at every point along the path generated by $\hat{\sigma}$. Thus, if the initial node $s_0 \epsilon \mathscr{D}(\epsilon)$ we say that *alternation is implementable* (in MPE), meaning that a tacit agreement to share the market through simple "turn-taking" could be implemented by a state–contingent, self-enforcing rule. Corollary 2 shows how the implementability of this norm depends on the size and distribution of capacities.

If demand is sufficiently inelastic, in that (4.3) holds, then $\mathscr{D}(\epsilon)$ contains any node

from which alternation never causes the stocks to differ by more than n', without regard to total capacity. So, for $\epsilon \leq \frac{\delta}{1+\delta}$, implementability depends only on symmetry. The higher the interest rate or longer the period length, the less elastic demand must be for this result to hold. With elastic demand (4.4) implies a trade—off between the aggregate size of capacity and the inequality of its distribution, in determining whether alternation is implementable. For $\epsilon > \frac{\delta}{1+\delta}$, $\phi(n,k) \geq \epsilon$ only if $\delta^n \geq \epsilon$ (i.e., $n \leq n'$) and k is not too large; the smaller is n, the larger can k be while maintaining $\phi(n,k) \geq \epsilon$. Thus, the greater is the total capacity, the more equally must it be held for alternation to be implementable.

Another implication of Theorem 2 is the Proposition below. Its proof examines the limit as $h \to 0$ of the set $h \cdot \mathscr{D}(\epsilon)$, and uses Theorem 1 to extend $\hat{\sigma}$ to a full MPE on S. This method foreshadows limit arguments developed more fully in section V of the paper.

Proposition 1: For any fixed $\epsilon \in (0,1)$, $r > 0$, and total capacity $\overline{R} \equiv R_A + R_B$, an MPE exists which achieves collusion whenever it is relevant, given that the periods are short enough and the initial holdings are sufficiently symmetric. Formally:

Let $\overline{R}$, ϵ, and r be given, and consider the vector of capacities $R \equiv (R_A, R_B)$, where $R_A + R_B = \overline{R}$. There exists $\overline{x} > 0$, $\overline{h} > 0$, and $\sigma \in \Sigma^*$ such that $\{s_t(s_0(h,R), \sigma)\} \subset M(\sigma)$ whenever $h < \overline{h}$, $|R_A - R_B| < \overline{x}$, and $Q^*(s_0(h,R)) = 1$.

Proof: See the appendix.

The proof is based upon the observation that, <u>when the periods are short and the firms equally placed, private incentives to conform with the alternating profile nearly coincide with "social" incentives to maximize industry revenues</u>: If each firm has $k \approx \overline{R}/2h$ units, then (4.4) implies that neither firm deviates from $\hat{\sigma}$ if $\epsilon \leq (\frac{\delta}{1+\delta})(1+\delta^{2k-1})$. On the other hand, coordination is required to maximize industry revenues if $\epsilon \leq \frac{1}{2}(1+\delta^{2k-1})$, by (4.1). But, as $h \to 0$ and $\delta \to 1$, these condition both require that $\overline{R} < -\ln(2\epsilon-1)/r$. Thus, the notions of "alternation" and "coordination" are intrinsically linked in this model.

D. MPE Extension of the Alternating Convention $\hat{\sigma}$

We now complete the program initiated in Theorem 2 by exhibiting a profile which is a perfect equilibrium extension of $\hat{\sigma}$ to the entire lattice. The extension profile takes the following general form on $S \backslash \mathscr{D}(\epsilon)$: The larger firm supplies one unit at every node, while the smaller firm supplies nothing if and only if, by so doing, the equilibrium path reaches $\mathscr{D}(\epsilon)$ within a given number of periods. The critical number of periods depends only on the smaller firm's endowment. The argument is based on a backward induction in two dimensions, as suggested by the discussion of section III. We require the following:

Definition 4.3: For strictly positive $s \in S$, and $0 < \epsilon < 1$, the "waiting time" $m(s)$ is given by:

(4.5)
$$m(s) \equiv \min\{m \in I \,|\, \forall s,\ 1 \le k \le k(s) : \phi(n(s)+1-m,k) \ge \epsilon\}, \text{ if } \phi(1,k(s)) \ge \epsilon, \text{ and}$$
$$m(s) \equiv +\infty, \text{ if } \phi(1,k(s)) < \epsilon.$$

Theorem 3: With any arbitrary sequence $\{m_k : k \ge 1\} \subset I$ associate the profile $\sigma \in \Sigma$ which extends $\hat{\sigma}$ as follows: On $\mathscr{D}(\epsilon)$, $\sigma(s) \equiv \hat{\sigma}(s)$. On $S \backslash \mathscr{D}(\epsilon)$, $\sigma_{j(s)}(s) \equiv 1$; $\sigma_{i(s)}(s) = 0$ if $1 \le m(s) \le m_{k(s)}$, otherwise $\sigma_{i(s)}(s) = 1$. For each $\epsilon \in (0,1)$ there exists a unique sequence $\{m_k^* : k \ge 1\}$ such that the associated profile σ^*, determined in this way, is an MPE.
Proof: See the appendix.

Comparing (4.4) and (4.5), in light of Corollary 2, one may see that $m(s)$ is the time that the smaller firm [i(s)] would have to wait before reaching $\mathscr{D}(\epsilon)$ if it does not produce, given that the larger firm [j(s)] supplies constantly. With every sequence $\{m_k\}$ we have associated this strategy for i on $S \backslash \mathscr{D}(\epsilon)$: "Wait for $\mathscr{D}(\epsilon)$ when holding k units if and only if you can get there within m_k periods; otherwise, produce." So $\{m_k^*\}$ may be interpreted as the parameters of an "optimal waiting policy" for i, given the assumed

behavior of j. The Theorem asserts that, given alternation on $\mathscr{D}(\epsilon)$, this "optimal waiting" by i together with constant production by j on $S \backslash \mathscr{D}(\epsilon)$ are mutual best responses.

This waiting strategy extends the set of nodes at which simultaneous production is avoided (and thus collusion is achieved) beyond the region of alternation, $\mathscr{D}(\epsilon)$, by giving precedence (and therefore, enhanced payoff) to the firm with greater capacity. An interesting interpretation of the MPE σ^* is that <u>the larger firm exploits its position so as to garner a greater than proportionate share of the collusive profits by credibly threatening to "produce all the time"</u> on $S \backslash \mathscr{D}(\epsilon)$, thereby <u>persuading the smaller firm that it is in its interest to forebear</u>, whenever $m(s) \leq m^*_{k(s)}$. To see the basis of this outcome consider a node s such that $k(s)=1$, and $n(s)=n>n'-1$. We noted above that $\sigma_{j(s)}(s)=1, \forall \sigma \epsilon \Sigma^*$, at such a node. Thus j has the dominant play (or, credible threat) of "1" at s; i knows this, and will wait as long as it can have the market to itself within n' periods. This reasoning, applied recursively, generates the full sequence of optimal waiting times, $\{m^*_k\}$.

When demand is inelastic in the sense of (4.3) these waiting times can be explicitly found. We give the derivation in the text to illustrate the conceptual basis of Theorem 3.

Corollary 3: If $\epsilon \leq \frac{\delta}{1+\delta}$ then the sequence $\{m^*_k\}$, $k \geq 1$, determining σ^* is given by:

$$(4.6) \qquad m^*_k = \text{Max}\{m \epsilon I \mid \delta^m \cdot (\frac{\delta}{1+\delta}) \cdot (1+\delta^{2k-1}) \geq \epsilon\}, \ k \epsilon I.$$

Proof: We use a nested induction to construct i's best response to $\{\sigma^*_j \equiv 1\}$ on $S \backslash \mathscr{D}(\epsilon)$. If $s \epsilon S \backslash \mathscr{D}(\epsilon)$ and $k(s)=1$, then it is obvious that i(s) either produces at s, earning $\beta\epsilon$, or waits for $\mathscr{D}(\epsilon)$, earning $\beta\delta^{m(s)}$. So i waits if and only if $m(s) \leq n'-1 \equiv m^*_1$, since, by the proof of Corollary 2, firm j produces under $\hat{\sigma}$ at the first point reached in $\mathscr{D}(\epsilon)$.

Inductively, assume that following σ^*_i is optimal for i(s) whenever $k(s) \leq k-1$, for $\{m^*_1, ..., m^*_{k-1}\}$ as given in (4.6), $k \geq 2$, and consider i's best play from $s \epsilon S \backslash \mathscr{D}(\epsilon)$ where $k(s)=k$. The proof rests on the fact that $m(s) \leq m^*_k$ implies $m(s-(1,1)) \leq m^*_{k-1}$ when $\epsilon \leq \frac{\delta}{1+\delta}$:

Corollary 2 implies that $m(s)=n(s)-N$, for $N\in\{n',n'-1\}$, and one easily verifies from (4.6) that $m_k^*\in\{m_{k-1}^*, m_{k-1}^*-1\}$, $k\geq2$, so $m(s-(1,1))=m(s)\leq m_k^*\leq m_{k-1}^*$. Hence, if firm i plays "1" at s, continuing optimally from $s-(1,1)$, it earns: $\beta\cdot[\epsilon+\delta^{m(s)+2}\cdot(1-\delta^{2k-2})/(1-\delta^2)]$. If it waits for $\mathscr{D}(\epsilon)$ it earns: $\delta^{m(s)+1}\cdot(1-\delta^{2k})/(1-\delta^2)$. Simple algebra shows waiting is better. Therefore, by induction, σ_i^* is optimal from all s such that $1\leq m(s)\leq m_{k(s)}^*$.

On the other hand, if $k(s)=k$ and $m(s)=m_k^*+1$, the same algebra now shows $i(s)$ does better to play "1" at s. Suppose this is true when $k(s)=k$ and $m_k^*<m(s)\leq m_k^*+m-1$, $m\geq2$, and consider the best play for $i(s)$ when $k(s)=k$ and $m(s)=m_k^*+m$. If it plays "1" at s and then "0" at $s-(1,1)$, it arrives at some node (say s') after two periods, having earned $\beta\epsilon$. If, alternatively, i plays "0" at s and then "1" at the successor node, it also arrives at s' after two periods, but now having earned only $\delta\beta\epsilon$ along the way. Clearly the former course dominates the latter. The induction hypothesis on k implies i will do no worse than the former course if it plays "1" at s, while the induction hypothesis on m implies i can do no better than the latter course if it plays "0". Therefore, induction on m, for fixed k, followed by induction on k, establishes that σ_i^* is optimal for i at any s for which $m(s)>m_{k(s)}^*$, completing the proof. $\square$

Finally, we note that the MPE σ^* obeys a kind of "inverse entropy" law: The market never evolves toward less coordinated play as capacities decline. For $s\in S\setminus\mathscr{D}(\epsilon)$, coordination is achieved at s if and only if $m(s)\leq m_{k(s)}^*$, in which case the ensuing path leads directly to $\mathscr{D}(\epsilon)$, with firm i forebearing until this region is reached, and alternating play thereafter. If $m(s-t\cdot(1,1))>m_{k(s)-t}^*$, $0\leq t\leq k(s)-1$, then both firms produce until i has exhausted its capacity. If $m(s)>m_{k(s)}^*$ but $m(s-t\cdot(1,1))=m_{k(s)-t}^*$, $1\leq t\leq k(s)-1$, the equilibrium path begins with both producing, but i withdraws at the node $s-t\cdot(1,1)$ and waits for $\mathscr{D}(\epsilon)$. So σ^* satisfies: $s\in M(\sigma^*)$ and $Q^*(s)=1$ only if $\{s_t(s,\sigma^*)\}\subset M(\sigma^*)$. That is, if coordination is achieved at any node, it is achieved at every subsequent node.

E. The Limits to Collusion in Perfect Equilibrium

We complete this discussion of tacit collusion by showing that it is possible to achieve the industry's revenue maximum in an MPE from a considerably larger set of nodes than $M(\sigma^*)$ when demand is inelastic. At the same time, our constructive methods show why collusion is not possible in any equilibrium when capacities are very unequally distributed. For the rest of this section we assume that $\epsilon \leq \frac{\delta}{1+\delta}$. Denote by M^* the set of nodes $s \in S$ from which there exists some MPE $\sigma \in \Sigma^*$ satisfying: $\{s_t(s,\sigma)\} \subset M(\sigma)$.

Any collusive path $\{s_t(s,\sigma)\}$ occasionally passes between the firms' "regions of precedence", so that a date t is reached when $\sigma(s_t) + \sigma(s_{t+1}) = (1,1)$. From Corollary 1 we know that $V_j(\sigma)(s_t) \leq \delta\beta(\frac{1-\epsilon}{1-\delta})$ where j is the firm for which $\sigma_j(s_t) = 0$. If the difference in holdings is great at s there will be no path which maximizes industry revenue and keeps the larger firm's payoff below this bound. But this means that the larger firm would not have been willing to forebear in order to arrive at s in any revenue maximizing MPE. So any such path arriving at s must have begun with the smaller firm waiting. There is a limit on how long a firm would be willing to wait, which may be determined by an inductive argument similar to that used in the proof of Theorem 3. A limit is ascertainable in this way on the degree of disparity in initial holdings consistent with an MPE generating a collusive path. We can then use Theorems 1 and 2, and the notion of alternation, to construct an MPE generating a collusive path from any node within these limits.

Ignoring Markov strategies momentarily, let us study those time paths of output which avoid simultaneous production everywhere without either firm's implied payoff becoming "too large". Call the path $\{q_t : 0 \leq t \leq \bar{t}\}$ *admissible from* s if (1) $q_t \in \{(0,1),(1,0)\}$, $\forall t$; (2) $q_0 + q_1 + ... + q_{\bar{t}} = s$; and (3) the value of the sequence $\{q_{t'} : t \leq t' \leq \bar{t}\}$, for either firm and any intermediate date t, $0 \leq t \leq \bar{t}$, does not exceed $\beta(\frac{1-\epsilon}{1-\delta})$. Consider the recursion:

(4.7)

(i) $\bar{n}_0 = n'$, and $\bar{n}_k = \bar{n}_{k-1} + c_k - 1$, $k \in I$, where for $k = 1, 2, \ldots,$

(ii) $c_k \equiv \text{Max}\{c \in I \mid \beta(\frac{1-\delta^c}{1-\delta}) + \delta^{1+c} \cdot V_{k-1} \leq \beta(\frac{1-\epsilon}{1-\delta})\}$, and

(iii) $V_k \equiv \beta(\frac{1-\delta^{c_k}}{1-\delta}) + \delta^{1+c_k} \cdot V_{k-1}$, with $V_0 \equiv \beta(\frac{1-\delta^{n'}}{1-\delta})$.

Then we have:

Lemma 3: An admissible path from s exists iff $n(s) \leq \bar{n}_{k(s)}$, where $\{\bar{n}_k\}$, $k \in I$, solves (4.7). Moreover, for $k \in I$, $\underline{c} \leq \bar{n}_k - \bar{n}_{k-1} + 1 \leq \underline{c} + 1$, where $\underline{c} \equiv \text{Max}\{c \in I \mid \delta^c(1-\delta)/(1-\delta^{c+1}) \geq \epsilon\} \geq 1$.

Proof: Consider the problem of finding the longest possible sequence of "1's" and "0's", given that exactly k "0's" are to be used in the sequence, such that no "tail" of the sequence ever has a present value exceeding $\beta(\frac{1-\epsilon}{1-\delta})$, where a "1" earns the value β, a "0" earns nothing, and the discount factor is δ. Let $\bar{n}_k + k$ be the maximal number of "1's" in this sequence, $k \in I$. With $\{\bar{n}_k\}$ so defined, it is obvious that an admissible path exists from s if and only if $n(s) \leq \bar{n}_{k(s)}$. The recursion (4.7) is the Bellman equation for this optimization problem: c_k is the number of "1's" which may be added to the sequence as the number of allowable "0's" increases from $k-1$ to k. So, $c_k = (k + \bar{n}_k) - (k - 1 + \bar{n}_{k-1})$, which is (i). V_k is the present value of a sequence at its maximal length, given k. So (ii) and (iii) follow from the Principle of Optimality and the definition of present value.

From (ii) and (iii): $\beta(\frac{1-\epsilon/\delta}{1-\delta}) < V_k \leq \beta(\frac{1-\epsilon}{1-\delta})$, $k \in I$. Substitution into (ii) yields:

(4.8) $\underline{c} \equiv \text{Max}\{c \in I \mid \delta^c(1-\delta)/(1-\delta^{c+1}) \geq \epsilon\} \leq c_k < \text{Min}\{c \in I \mid \delta^c(1-\delta)/(1-\delta^c) < \epsilon\} \equiv \bar{c}$, $k \geq 1$.

Since $\frac{\delta}{1+\delta} \geq \epsilon$, $\underline{c} \geq 1$; while the definitions above imply $\bar{c} \leq \underline{c} + 2$, and hence $\underline{c} \leq c_k \leq \underline{c} + 1$. Moreover, if $\delta^{\underline{c}}(1-\delta)/(1-\delta^{\underline{c}+1}) \geq \epsilon > \delta^{\underline{c}+1}(1-\delta)/(1-\delta^{\underline{c}+1})$, then $c_k = \underline{c}$, $\forall k \geq 1$. □

Define $\bar{N} \equiv \{s \in S \mid k(s) = 0 \text{ or } n(s) \geq \bar{n}_{k(s)} - c_{k(s)}\}$, where $\{\bar{n}_k\}$ and $\{c_k\}$, $k \geq 1$, are as given in (4.7). Notice that if $n(s) > \bar{n}_{k(s)} - c_{k(s)}$ then, by Lemma 3 and Corollary 1, no col

lusive MPE path is possible in which the smaller firm produces at s. Thus, any MPE path from $s \in \bar{N}$ attaining an industry revenue maximum involves firm $i(s)$ waiting at least $n(s) - [\bar{n}_{k(s)} - c_{k(s)}]$ periods before producing. Once it has produced on a collusive MPE path at some node $s' \in \bar{N}$ for which $n(s') = \bar{n}_{k(s')} - c_{k(s')}$, the ensuing path, necessarily admissible, must follow the pattern of production implied by (4.7): c_{k-t} units supplied by the larger firm, followed by one unit from the smaller firm, for $1 \leq t \leq k-1$. Thus, although admissibility is an "open–loop" concept, it provides a powerful tool for assessing the limits to collusion in MPE. For in the Appendix we prove that if an admissible path exists from $s \in S$, or if $i(s)$ would wait from s to reach the only admissible path in $\bar{N}$, then $s \in M^*$.

Theorem 4: Given any sequence $\{m_k : k \geq 1\} \subset I$ define the profile $\tilde{\sigma}$ on $\bar{N}$ as follows:

    (a) If $n(s) = \bar{n}_{k(s)} - c_{k(s)}$ then $\tilde{\sigma}_{j(s)} = 0$ and $\tilde{\sigma}_{i(s)} = 1$;

    (b) If $k(s) = 0$, or $m_{k(s)} + \bar{n}_{k(s)} \geq n(s) > \bar{n}_{k(s)} - c_{k(s)}$, then $\tilde{\sigma}_{j(s)} = 1$ and $\tilde{\sigma}_{i(s)} = 0$;

    (c) If $n(s) > \bar{n}_{k(s)} + m_{k(s)}$ then $\tilde{\sigma}(s) = (1,1)$.

Then: (i) $\exists \{\overline{m}_k : k \geq 1\}$ for which $\tilde{\sigma} \in \Sigma^*_{|\bar{N}}$; moreover, (ii) $\{s \in S \mid n(s) \leq \bar{n}_{k(s)} + \overline{m}_{k(s)}\} \subset M^*$.

Proof: See the Appendix.

    The kinship of Theorems 3 and 4 should be apparent. Theorem 3 constructs the MPE σ^* from the restricted equilibrium $\hat{\sigma}$ on $\mathscr{D}(\epsilon) \subset S$, showing that constant supply by the larger firm, and optimal waiting by the smaller one, are mutual best responses on $S \backslash \mathscr{D}(\epsilon)$. Theorem 4 constructs the restricted MPE $\tilde{\sigma}$ on $\bar{N}$ by specifying behavior on the "boundary" of $\bar{N}$ to coincide with the unique admissible path in that set, and then showing that constant supply by j together with optimal waiting by i are mutual best responses, given this specified boundary behavior. Since $\tilde{\sigma}$ is a restricted MPE on $\bar{N}$, there exists some $\overline{\sigma} \in \Sigma^*$ such that $\overline{\sigma} \equiv \tilde{\sigma}$ on $\bar{N}$. The second part of Theorem 4 is based on the fact that, for any $s \notin \bar{N}$, we can choose $\overline{\sigma}$ so that the path $\{s_t(s, \overline{\sigma})\}$ implies alternating play on $S \backslash \bar{N}$, and hence $\{s_t(s, \overline{\sigma})\} \subset M(\overline{\sigma})$. We conjecture, but have as yet not proven,

that $\tilde{\sigma}$ completely characterizes M^* in the sense that if $n(s) > \bar{n}_{k(s)} + \overline{m}_{k(s)}$, then $s \notin M^*$.

Theorem 4(ii) states that the existence of an admissible path from a given node implies some MPE attains an industry revenues maximum from that node. So Lemmas 2 and 3 together show clearly why coordination via non–cooperative play is more difficult to achieve when firms precommit to paths of actions than when they use state–contingent rules, in this dynamic duopoly with indivisible production. With open–loop strategies, the possibility of coordination turns on the ability to find a path for which the discounted value of the high price on each firm's last supply date is no less than the discounted value of the low price on his first date of forebearance. With Markov play what matters is whether the right to produce can be allocated so that neither firm has to wait when the value of continuing, on a prorated per period basis, exceeds the fall in price induced by a loss of coordination. Comparing (4.2) and (4.7) reveals that the former is the more severe constraint.

Theorem 4 suggests that the concepts of "alternation", "forebearance", and "optimal waiting" are sufficient to fully describe how the firms might tacitly coordinate their actions in this setting. As the proof of the Theorem reveals, the extreme admissible path implied by (4.7) constitutes a "boundary" separating the "interior" of $\overline{N}$ from the rest of S. From any $s \in S \backslash \overline{N}$ the firms agree to alternate until reaching this boundary; from any $s \in \overline{N}$ not on the boundary, the smaller firm waits for the boundary if it is in its interest to do so. Play along the boundary gives the larger firm sole access to the market roughly the fraction $1 - \epsilon$ of the time [see (4.8)]. This tacit agreement is self–enforcing, because each firm knows it to be in the other's interest to conform from any node in the state space.

V. Depicting Discrete Equilibrium Dynamics in a Continuous State Space

Consider a continuous formulation of this game, without the assumption that output rates are fixed for intervals of length h. The state space is all capacity pairs, $\mathbb{R}_+^2$. Each firm's (Markov) strategy maps capacities into production rates: $\hat{q}_i : \mathbb{R}_+^2 \to \{0,1\}$, $i = A, B$. The profile $\hat{q} = (\hat{q}_A, \hat{q}_B)$ partitions the state space into regions of production, $\hat{q}_i^{-1}(1)$, and

forebearance, $\hat{q}_i^{-1}(0)$, for each firm, and induces a vector field determining the evolution of the system over time: $\frac{dR}{d\tau} = -\hat{q}(R_\tau)$, R_0 given. For $R = (R_A, R_B)$, the associated phase diagram shows motion "due South" on $\hat{q}_A^{-1}(0) \cap \hat{q}_B^{-1}(1)$, "due West" on $\hat{q}_A^{-1}(1) \cap \hat{q}_B^{-1}(0)$, and "Southwest" on $\hat{q}_A^{-1}(1) \cap \hat{q}_B^{-1}(1)$. Payoffs from any node are given in the obvious way, as the discounted integral of instantaneous revenue flows. An equilibrium consists of a strategy pair from which neither firm can affect a unilateral improvement in its payoff from any node, where strategies are restricted so that the differential equation describing the evolution of the system always has a solution. Any such restriction will imply that the sets $\hat{q}_i^{-1}(1)$ have a relatively simple topological structure.

Now, let $\{\sigma^h : h > 0\}$ be a family of MPE of the discrete game, indexed by the period length h. For h fixed, σ^h will cause the vector of capacities, R_τ, to evolve over time in the following manner: Given R_0, a unique intitial state $s_0(h, R_0) \approx h^{-1} \cdot R_0$ is determined, as is the path $\{s_t(s_0(h, R_0), \sigma^h)\}$. So,

$$R_{\tau+h} - R_\tau = -h \cdot \sigma^h(s_t(s_0(h, R_0), \sigma^h)) = -h \cdot \sigma^h(s_0(h, R_\tau)), \quad \forall \tau = h \cdot t, \ t \in I.$$

By analogy with the continuous case, $\{\sigma^h : h > 0\}$ defines a vector field on $\mathbb{R}_+^2$ given by:

$$\frac{dR}{d\tau} = -\hat{q}(R_\tau) \equiv -\lim_{h \to 0} \{\sigma^h(s_0(h, R_\tau))\},$$

if the limit exists. The implied phase diagram, showing the motion over time of $(R_\tau : \tau \geq 0)$, would mimic closely the path of capacities and the payoffs generated by σ^h, for "small" h.

The MPE σ^* defined in Theorem 3 treats h parametrically, so it actually defines a family of MPE, $\{\sigma^{*h} : h > 0\}$, as envisioned above. But σ^{*h} is an extension of $\hat{\sigma}^h$, the alternating profile, and $\{\hat{\sigma}^h(s_0(h, R))\}$ has no limit, as $h \to 0$: Alternation leads to "chattering" for infinitesmal h. The continuous approach outlined above could not produce behavior in equilibrium like that which emerges in the equilibria we have studied for the discrete game, with h "small". This is not necessarily a criticism of the discrete approach; one might as readily conclude that the continuous formulation is flawed, to the extent that it fails to capture the strategic possibilities inherent in a convention of "turn–taking".

Moreover, this "chattering" in the limit need not prevent a phase diagrammatic depiction of equilibrium dynamics; we can simply define: $(\frac{1}{2},\frac{1}{2}) \equiv \lim\{\hat{\sigma}^h(s_0(h,R))\}$, as $h\to0$, $\forall R\in\mathbb{R}^2_{++}$.

With this definition in mind, note that the equilibrium σ^* defined in Theorem 3 partitions the state space S into three mutually exclusive and collectively exhaustive regions: a region of alternation, $\mathcal{D}$; a region of forebearance by the smaller firm, $\mathcal{F}$; and a region of simultaneous production, $\mathcal{S}$. Drop the explicit dependence of $\mathcal{D}$ on ϵ. Clearly, $\mathcal{F}=\{s\in S\,|\,0<m(s)\le m^*_{k(s)}\}$ and $\mathcal{S}=\{s\in S\,|\,m(s)>m^*_{k(s)}\}$. The dynamics generated by σ^*, for "small" h, satisfy: $\frac{dR}{d\tau}=-(\frac{1}{2},\frac{1}{2})$, if $s_0(h,R)\in\mathcal{D}$ as $h\to0$; $\frac{dR}{d\tau}=-(1,1)$, if $s_0(h,R)\in\mathcal{S}$ as $h\to0$; and, $\frac{dR}{d\tau}=-(1,0)\ [-(0,1)]$, if $s_0(h,R)\in\mathcal{F}$ as $h\to0$, for $j(s_0(h,R))=A\ [B]$. The results in Section IV (and their proofs) can be used to solve for the regions of $\mathbb{R}^2_+$ in which these various motions obtain, even when explicit formulae cannot be found in the discrete case.

Abusing notation, write $R\in\mathcal{D}\ [\mathcal{F},\ \mathcal{S}]$ if $\exists\bar{h}>0$ such that, $\forall h\le\bar{h}$, $s_0(h,R)\in\mathcal{D}\ [\mathcal{F},\ \mathcal{S}]$. Denote $x\equiv|R_A-R_B|$, $y\equiv\text{Min}\{R_A,R_B\}$ and $z\equiv\text{Max}\{R_A,R_B\}$; then $h\to0\Rightarrow h\cdot n(s_0(h,R))\to x$ and $h\cdot k(s_0(h,R))\to y$. Define $w(x,y)\equiv\lim_{h\to0}\{h\cdot m(s_0(h,R))\}$, and $w^*(y)\equiv\lim_{h\to0}\{h\cdot m^*_{k(s_0(h,R))}\}$.

Proposition 2: The limiting dynamics implied by the MPE (family) σ^* on the continuous state space $\mathbb{R}^2_+=(\mathcal{D}\cup\mathcal{F}\cup\mathcal{S})$ are as described above, where:

$$R\in\mathcal{D} \text{ if } e^{-rx}\cdot e^{-2ry'}+\frac{1}{2}\cdot(1-e^{-2ry'}) > \epsilon, \forall y'\in[0,y];$$

$$R\in\mathcal{F} \text{ if } 0 < w(x,y) < w^*(y); \text{ and, } R\in\mathcal{S} \text{ if } w(x,y) > w^*(y).$$

Moreover, $w(x,y) = x+\frac{1}{r}\ln(\epsilon)$, $\epsilon\le\frac{1}{2}$; $w(x,y) = x+\frac{1}{r}\cdot\ln[\frac{1}{2}+(\epsilon-\frac{1}{2})e^{2ry}]$, $\epsilon>\frac{1}{2}$; and:

(5.1)
$$w^*(y) = \frac{1}{r}\cdot\ln[(1+e^{-2ry})/2\epsilon], \quad \epsilon\le\frac{1}{2};$$

(5.2)
$$w^*(y) = \text{Max}\{0, \frac{1}{r}\cdot\ln\left[\frac{u(y)+[2-u(y)]e^{-2ry}}{2\epsilon}\right]\}, \quad \epsilon>\frac{1}{2},$$

for $u(y) \equiv \left[\dfrac{1-(2\epsilon-1)e^{2ry}}{1+(2\epsilon-1)e^{2ry}}\right]$, $0 \leq y \leq \bar{y} \equiv -\ln(2\epsilon-1)/2r$. When $\epsilon > \frac{1}{2}$, $w^*(y)=0$, $\hat{y} \leq y \leq \bar{y}$, for $\hat{y} \in (0,\bar{y})$.

Proof: The argument is a straightforward consequence of taking limits in (4.4), (4.5) and (4.6), except for the derivation of (5.2), which is in the Appendix. By Corollary 2, $R \in \mathscr{D}$ if $\exists \bar{h} > 0$ such that, $\forall h \leq \bar{h}$: $\phi(n(s_0(h,R))+1,k) \geq \epsilon$, $\forall k \in I$, $k \leq k(s_0(h,R))$. Then (4.4) implies:

$$\phi[n(s_0(h,R))+1,\ k(s_0(h,R))] \to e^{-rx} \cdot e^{-2ry} + \frac{1}{2} \cdot (1-e^{-2ry}), \text{ as } h \to 0.$$

Moreover, from (4.5) one sees that $m(s_0(h,R))$ must "solve": $\phi(n-m+1,k) \approx \epsilon$. Equating the limit above with ϵ, substituting $(x-w)$ for x, and solving for w, gives $w(x,y)$. Letting $h \to 0$ in (4.6) gives: $\epsilon \approx \delta^m(\frac{\delta}{1+\delta})(1+\delta^{2k-1}) \to \frac{1}{2} \cdot e^{-rw^*} \cdot (1+e^{-2ry})$, as $h \to 0$, which is (5.1). The derivation of (5.2) is more complex since no explicit representation of $\{m_k^*\}$ is available when $\epsilon > \frac{\delta}{1+\delta}$; equation (A.2) in the proof of Theorem 3 must be utilized. □

It follows from (5.1) that, when $\epsilon \leq \frac{1}{2}$, a coordination failure occurs under the MPE σ^* for "small" h if $z = x+y > \frac{1}{r} \cdot \ln\left[\dfrac{e^{ry}+e^{-ry}}{2\epsilon^2}\right]$. The comparable limit in (4.2) implies that a coordination failure must occur in any open-loop equilibrium if $z > \frac{1}{r} \cdot \ln\left[\dfrac{e^{-2ry}}{\epsilon^2}\right]$. This gives a clear measure of the superiority of closed-loop play for achieving coordination. A comparable measure can be derived from (5.2), which shows that the advantage of state-contingent play vanishes as $\epsilon \to 1$. Observe that the expressions in (5.1) and (5.2) converge as $\epsilon \to 1/2$, so equilibrium behavior is continuous in the elasticity of demand. It is easy to compute firms' payoffs for various endowment configurations using the results of Proposition 2, and so to verify our assertion that the larger firm garners a greater than proportionate share of industry earnings from nodes $R \in \mathscr{F}$ where $w(x,y) \approx w^*(y)$. Figures 4A and 4B depict the limit phase diagrams for the MPE σ^* in the respective cases $\epsilon \leq \frac{1}{2}$ and $\epsilon > \frac{1}{2}$.

The power of this limit argument is made clear by its application to the restricted MPE $\tilde{\sigma}$ defined in Theorem 4. For there the discrete game is very difficult to analyze,

Figure 4A

$z = y - \ln(\varepsilon)/r$

$\dot{R} = -(\frac{1}{2}, \frac{1}{2})$

$\dot{R} = -(1,1)$

$\dot{R} = -(1,0)$

$z = \ln\left(\dfrac{e^{ry} + e^{-ry}}{2\varepsilon^2}\right)$

$45°$ $45°$

Figure 4B

$\bar{y}$

$\dot{R} = -(1,1)$

$\hat{y}$

$\dot{R} = -(\frac{1}{2}, \frac{1}{2})$ $\dot{R} = -(1,0)$

$0 = w(x, \hat{y})$ $w(x,y) = w^*(y)$

45
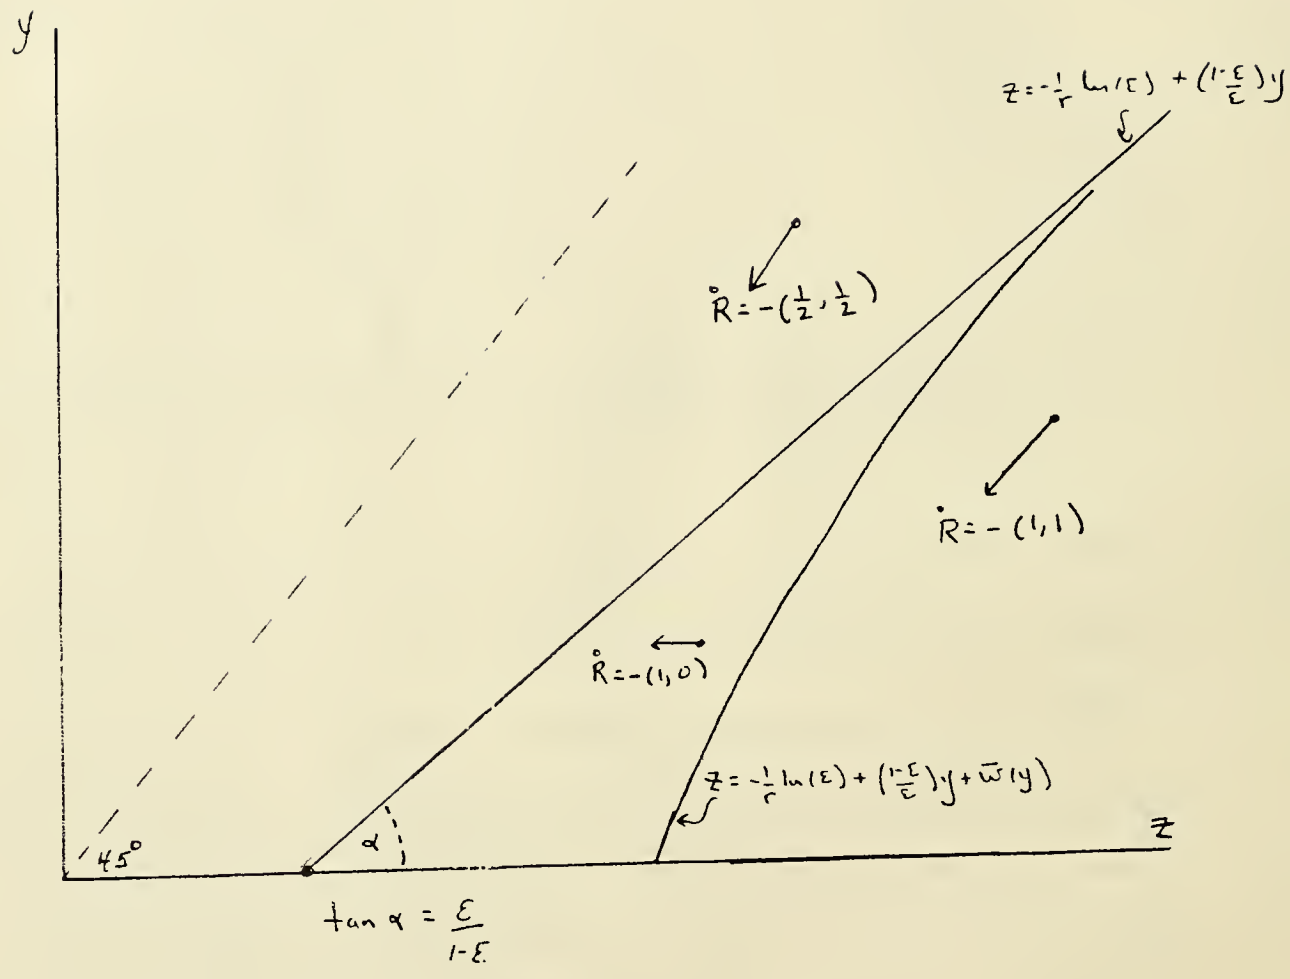
Figure 4c

but its continuous analogue is transparent. Lemma 3, and equations (4.7) and (4.8), can be employed to see that $s_0(h,R)\epsilon\overline{N}$, for h sufficiently small, if $z > -\ln(\epsilon)/r + (\frac{1-\epsilon}{\epsilon})\cdot y$. The extreme admissible path defining the boundary of $\overline{N}$ "chatters" with infinitesmal h, since (by (4,8)) $c_k\epsilon\{\underline{c}, \underline{c}+1\}$, where $\underline{c}=\text{Max}\{c\epsilon I | c \le \frac{1-\epsilon}{\epsilon}\}$. But as $h \to 0$ this boundary becomes a line with slope (relative to the larger firm's axis) $\frac{\epsilon}{1-\epsilon}$, along which the larger firm produces on average the fraction $1-\epsilon$ of the time. (See Figure 4C.) The proof of Theorem 4 shows that one can extend $\tilde{\sigma}$ in such a way that, "above" this line, the firms agree to take alternate turns. The profile $\tilde{\sigma}$ requires that "below" this line the larger firms produces constantly and the smaller firm waits optimally. Optimal waiting is equivalent to the efficient use of an option to go to the boundary of $\overline{N}$, taking the payoff implied by $\tilde{\sigma}$ there.

Let us see how the limit of the waiting times, $\overline{w}(y) \equiv \lim_{h \to 0}\{h\cdot\overline{m}_{k(s_0(h,R))}\}$, and hence the behavior implied by $\tilde{\sigma}$, may be derived. From equation (A.3) in the Appendix one can verify that the problem facing the smaller firm at any node s for which $n(s) > \bar{n}_{k(s)}$ is to decide how long to earn $\beta\epsilon$ by producing each period, before waiting to move to the extreme admissible path where the firm can earn β each time it produces, but only gets to produce the fraction ϵ of the time. (Because of discounting, earning β for duration ϵ, and nothing for duration $1-\epsilon$ is better than earning $\beta\epsilon$ continuously, so waiting might pay.) For each unit of time that the firm produces, the wait increases by roughly $\frac{1-2\epsilon}{\epsilon}$ units of time, since the boundary of $\overline{N}$ has slope $\frac{1-\epsilon}{\epsilon}$ while the path along which both firms produce has slope 1. Passing to the limit, if the smaller firm has y units of stock, if the wait for the boundary of $\overline{N}$ is initially w, and if the firm produces for t units of time before beginning to wait, then its earnings are:

(5.3) $\qquad \Pi(y,t,w) = \frac{\epsilon}{r}\cdot\{[1-e^{-rt}] + e^{r(t-w)}[e^{-rt/\epsilon}-e^{-ry/\epsilon}]\}$, $\quad y>0$, $w\ge0$, and $0\le t\le y$.

Accordingly:

(5.4) $\qquad\qquad \overline{w}(y) \equiv \text{Max}\{w\ge0 | \Pi(y,0,w) \ge \underset{0 \le t \le y}{\text{Max}}\ \Pi(y,t,w)\}.$

One may verify from (5.3) that $\partial^2 \Pi / \partial w \, \partial t > 0$, so that $\Pi(y,0,w) \geq \Pi(y,t,w)$, $\forall t \in [0,y]$, if and only if $w \leq \overline{w}(y)$. Notice as well that Π is not generally a concave function of t. This fact has the implication that the equilibrium path under $\tilde{\sigma}$ may not be continuous in the initial conditions, since the argmax in (5.4) need not be a continuous function of w and y. [E.g, this continuity is crucial to the derivation of (5.2).] Thus, it may be desirable for the smaller firm to wait to reach the boundary of N from some node, and yet a slight decrease in endowment would cause its strategy to switch to producing for a protracted period. A number of interesting results along these lines can be derived with remarkable ease from (5.3) and (5.4), as exemplified by the following:

Proposition 3: Assume $\epsilon \leq 1/3$. Then $\overline{w}(y) \equiv \frac{1}{r} \cdot \ln\left[\frac{1 - e^{-ry/\epsilon}}{1 - e^{-ry}} \right]$. Some MPE extension of $\tilde{\sigma}$ attains the high price continuously if $z < -\ln(\epsilon)/r + (1-\epsilon)y/\epsilon + \overline{w}(y)$, for h sufficiently small. If the inequality is reversed, then the low price obtains at every date until the smaller firm has exhausted, under any MPE extension of $\tilde{\sigma}$, when h is small enough.

Proof: Using subscripts to denote partial derivatives, (5.3) implies:
$$r^{-1}\Pi_t(y,t,w) = e^{-rt} - e^{r(t-w)} \cdot [(\frac{1-\epsilon}{\epsilon})e^{-rt/\epsilon} + e^{-ry/\epsilon}]; \text{ while,}$$
$$r^{-2}\Pi_{tt}(y,t,w) = -e^{-rt} + e^{r(t-w)} \cdot [(\frac{1-\epsilon}{\epsilon})^2 e^{-rt/\epsilon} - e^{-ry/\epsilon}].$$
Thus: $[r^{-1}\Pi_t + r^{-2}\Pi_{tt}] = e^{r(t-w)} \cdot [(\frac{1-\epsilon}{\epsilon})(\frac{1-2\epsilon}{\epsilon})e^{-rt/\epsilon} - 2e^{-ry/\epsilon}] \geq 0$, $0 \leq t \leq y$, $\forall w \geq 0$, since $\epsilon \leq 1/3$. So $\Pi_t < 0 \Rightarrow \Pi_{tt} > 0$, and $\Pi_{tt} < 0 \Rightarrow \Pi_t > 0$. In particular, the function $\Pi(y,t,w)$ never attains an interior maximum with respect to t on the interval $[0,y]$. Thus, since $\Pi_{tw} > 0$, for any $y > 0$ there exists $\overline{w}(y)$ such that $\Pi(y,0,\overline{w}) \geq \Pi(y,y,\overline{w})$ iff $y \leq \overline{w}$. $\square$

VI. Conclusion

In the model of dynamic duopoly studied here firms regulate an "on–off" production

process, depleting their respective supplies of an non—replenishable factor of production at a constant rate whenever their process is "on". When both are producing the price is low. If only one of them produces an any moment the price is high. Each of them knows, at any moment, the remaining capacity to produce of the other. Two notions of equilibrium are investigated. In open—loop Nash Equilibrium the firms commit at the initial date to time paths of production which are mutual best responses. In Markov Perfect Equilibrium, the firms specify state—contingent plans of actions which are mutual best responses in every contingency. The capacity of the firms, in these alternative strategic settings, to achieve coordination of their actions via non—cooperative play has been systematically investigated.

By assuming firms can change their actions only on exogenously given, equally spaced moments in time, we convert the problem to a discrete time game which, due to the dichotomous production structure, can be explicitly solved. The possibility for coordination to be attained via non—cooperative play, in either strategic setting, is characterized (in Lemmas 2 and 3) by two relatively simple integer programming problems. These characterizations allow us to derive the main comparative results of the paper: that the ability to precommit undermines the possibility of coordination; and that state—contingent strategies are more effective at coordinating play when firms are symmetrically placed.

The analysis emphasizes a difference between discrete and continuous time dynamic games recently stressed by Simon and Stinchcombe (1989). The extensive form analyzed in the discrete game, though restricted to Markov strategies, reveals a rich structure largely hidden from view in a continuous time treatment. The idea of "precedence" — firms' making their play contingent on "who did what last period" — is fundamental to achieving maximal coordination in this environment. Of course, this finding is related to the "on—off" process of production which has been assumed here, but the point would seem to be of more general import. The results in section V on the limit behavior of our Markov Equilibria suggest the possibility of profitably using this discrete approach in other settings where continuous time formulations have been reflexively adopted.

Appendix

Proof of Theorem 1:

Define $S^k \equiv \{(a,b) \in S \mid a+b \leq k\}$, $k \in I$. The sets $\{S^k\}$ are nested, cover S, and are closed to deviations from any profile $\sigma \in \Sigma$. Therefore $\Sigma^*_{|S^{k+1}} \subset \Sigma^*_{|S^k}$, $k \geq 1$, and hence

$\Sigma^* = \bigcap \{\Sigma^*_{|S^k} : k \geq 1\}$. Now, clearly, $\Sigma^*_{|S^1} = \Sigma$. For fixed $k \geq 1$, for any $\sigma \in \Sigma^*_{|S^k}$, and for every

$s \in (S^{k+1} \backslash S^k)$, one of the alternatives (3.2)(a)–(c) holds at s. Moreover, for such k and s, the continuation payoffs in (3.2)(a)–(c) do not depend on values taken by σ on $S \backslash S^k$.

Therefore, given $W \subset S$ and $\bar{\sigma} \in \Sigma^*_{|W}$, the following inductive procedure constructing the profile $\sigma^* \in \Sigma$ is well defined. For successive values of $k=1, 2, ...$: (1) At each node $s \in (S^{k+1} \backslash S^k) \cap W$, set $\sigma^*(s) = \bar{\sigma}(s)$. (2) At each $s \in S^{k+1} \backslash (S^k \cup W)$, assign a value to σ^* consistent with the Bellman Inequalities there. Now $\bar{\sigma} \in \Sigma^*_{|W}$ implies $\bar{\sigma} \in \Sigma^*_{|S^k \cap W}$, $k \geq 1$. Thus

Lemma 1 implies that $\sigma^* \in \Sigma^*_{|S^k}$ at each step k. Therefore $\sigma^* \in \Sigma^*$.   □

Proof of Corollary 2:

Let $s \in S$ be strictly positive, with $n=n(s)$, $k=k(s)$, $i=i(s)$ and $j=j(s)$. By the definition of the alternating profile we have that:

(i) $V_j(\hat{\sigma})(s) = \delta\beta \cdot [(1-\delta^{2k-2})/(1-\delta^2) + \delta^{2k-2} \cdot (1-\delta^{n+1})/(1-\delta)]$, if $\hat{\sigma}_j(s)=0$; and,

(ii) $V_i(\hat{\sigma})(s) = \delta\beta(1-\delta^{2k})/(1-\delta^2)$, if $\hat{\sigma}_i(s)=0$.

Now, by Definition 4.1, $s \in \mathcal{D}(\epsilon)$ if, for every s' a strictly positive successor of s under $\hat{\sigma}$, $V_j(\hat{\sigma})(s') \leq \delta(\frac{1-\epsilon}{1-\delta})$ when $\hat{\sigma}_j(s')=0$, and $V_i(\hat{\sigma})(s') \leq \delta(\frac{1-\epsilon}{1-\delta})$ when $\hat{\sigma}_i(s')=0$. Notice that

the RHS of (i) exceeds the RHS of (ii), for any $n \geq 0$, $k \geq 1$. Therefore, when $\hat{\sigma}_j(s)=0$ then

$s \in \mathcal{D}(\epsilon)$ if: (1) [RHS (i)]$\leq \delta\beta(\frac{1-\epsilon}{1-\delta})$ for $1 \leq k \leq k(s)$ and $n=n(s)$; and, (2) [RHS (ii)]$\leq \delta\beta(\frac{1-\epsilon}{1-\delta})$ for

$1 \leq k \leq k(s)-1$. When $\hat{\sigma}_i(s)=0$ then $s \in \mathcal{D}(\epsilon)$ if: (1) [RHS (i)]$\leq \delta\beta(\frac{1-\epsilon}{1-\delta})$, for $1 \leq k \leq k(s)$ and

$n=n(s)-1$; and, (2) [RHS (ii)]$\leq \delta\beta(\frac{1-\epsilon}{1-\delta})$ for $1 \leq k \leq k(s)$. Condition (2) is redundant in each

case if $n \geq 1$. Given (4.4), condition (1) is equivalent to: $\phi(n(s),k) \geq \epsilon$, $1 \leq k \leq k(s)$, if $\hat{\sigma}_i(s)=0$;

and, $\phi(n(s)+1,k) \geq \epsilon$, $1 \leq k \leq k(s)$, if $\hat{\sigma}_j(s)=0$. But if (4.3) holds then $\phi(n,k) \geq \epsilon$, $\forall k \geq 1$, if and

only if $n \leq n'$. This completes the proof.   □

Proof of Proposition 1:

Suppose $\epsilon < 1/2$. Let $R=(R_A, R_B)$ be some vector of initial stocks such that

$R_A + R_B = \bar{R}$. Given any $h > 0$, Corollary 2 implies $s_0(h, R) \in \mathscr{D}(\epsilon)$, $\forall \bar{R} > 0$, if:

(i) $|R_A - R_E| < -\ln(\epsilon)/r$, and

(ii) $h < \ln(\frac{1-\epsilon}{\epsilon})/r$.

Condition (i) assures that $n(s_0(h)) \leq n'(h) - 1$ (neglecting rounding error, which vanishes for small h), while condition (ii) gaurantees that (4.3) holds. Theorems 1 and 2 together imply there exists an MPE profile σ coinciding with $\hat{\sigma}$ on $\mathscr{D}(\epsilon)$. Since $s_0 \in \mathscr{D}(\epsilon)$, from the definition of $\hat{\sigma}$, we conclude that $\{s_t(s_0, \sigma)\} \subset M(\sigma)$.

Suppose $\epsilon \geq 1/2$, and continue to denote by σ an MPE satisfying $\sigma \equiv \hat{\sigma}$ on $\mathscr{D}(\epsilon)$. Let $R = (R_A, R_B)$, with $R_A + R_B = \bar{R}$, and define $|R_A - R_B| \equiv x$, $\text{Min}\{R_A, R_B\} \equiv y$. Fix $h > 0$ and let $s = s_0(h, R)$. It follows from (4.4) that $s \in \mathscr{D}(\epsilon)$ when $\hat{\sigma}_{j(s)}(s) = 0$ if and only if:

(A.1)
$$\delta(h)^{2k(s)-1} \geq \frac{\epsilon \cdot [1 + \delta(h)] - \delta(h)}{\delta(h)^{n(s)}[1 + \delta(h)] - 1}.$$

Take the limit in (A.1) as $h \to 0$, using the definitions of s, $n(\cdot)$, and $k(\cdot)$, to see that $s_0(h, R) \in \mathscr{D}(\epsilon)$ for h sufficiently small if:

(iii) $e^{-2ry} > (2\epsilon - 1)/(2e^{-rx} - 1)$.

As long as:

(iv) $\bar{R} < -\ln(2\epsilon - 1)/r$,

it will be possible to select $x > 0$ and $y > 0$ such that $x + 2y = \bar{R}$, and (iii) holds. Now, by the analysis leading to (4.1) we know that $Q^*(s_0) = 1$ iff: $\bar{n}(s) \geq n(s) + 2k(s) \approx \bar{R}$. But (4.1) implies $h \cdot \bar{n}(s_0(h, R)) \to -\ln(2\epsilon - 1)/r$, as $h \to 0$. Thus (iv) is satisfied whenever $\bar{R}$ is below the threshold at which collusion becomes desirable. For such $\bar{R}$ it is always possible to divide the total endowment such that (iii) holds. Thus, for h is small enough we can find $R = (R_A, R_B)$, $R_A + R_B = \bar{R}$, such that $\{s_t(s_0(h, R), \sigma)\} \subset M(\sigma)$ whenever $Q^*(s_0(h, R)) = 1$. □

Proof of Theorem 3:
We have to show that the profile τ^* is well defined, and that neither firm will deviate from it at any node $s \in S \setminus \mathscr{D}(\epsilon)$. This amounts to showing two things: (1) that there exists a unique sequence of integers $\{m_k^*\}$ such that on $S \setminus \mathscr{D}(\epsilon)$ the strategy "wait for $\mathscr{D}(\epsilon)$ from s if you can get there within $m_{k(s)}^*$ periods" is a best response for $i(s)$ to the strategy "always produce" for $j(s)$; and, (2) the strategy "always produce" is a best response for $j(s)$ to the strategy given in (1) for $i(s)$. The proof of Corollary 3 establishes (1) in the case $\epsilon \leq \delta/(1+\delta)$, when an explicit formula for $\{m_k^*\}$ can be found. The nested induction employed there may be extended to the general case $\epsilon \in (0,1)$ as follows.
Consider $S(k) \equiv \{s \in S \setminus \mathscr{D}(\epsilon) | k(s) \leq k\}$, $k \geq 1$. Given any integers $\{m_1, ..., m_k\}$, the

profile defined in the Theorem, and thus each firm's value function, is uniquely determined at every node in $S(k)$. Therefore, the following induction uniquely defines a sequence of integers $\{m_k^*\}$, $k\geq1$, and so, by the definition in the Theorem, a profile σ^*:

(A.2) $m_1^* \equiv n'-1$; and for $k\geq2$,

$$m_k^* \equiv \operatorname*{Max}_{s \in S} \{m(s)\} \text{ s.t. } k(s)=k, \text{ and } \delta^{m(s)+1}\cdot[\frac{1-\delta^{2k}}{1-\delta^2}] \geq \beta\epsilon+\delta V_{i(s)}(\sigma^*)(s-(1,1)),$$

where $m_k^*\equiv0$ if the inequality always fails. This definition is meaningful because at each $k\geq2$ the profile σ^* and values $V_i(\sigma^*)(\cdot)$, $i=A,B$, are already defined on $S(k-1)\cup\mathscr{D}(\epsilon)$; and, $k(s)=k$ implies $s-(1,1)\in S(k-1)\cup\mathscr{D}(\epsilon)$. The integers m_k^* are defined by the requirement that waiting to reach an alternating path is dominated for $i(s)$ by producing whenever $m(s)>m_k^*$. What we have to show is that waiting for one period at s is optimal for $i(s)$ if and only if $m(s)\leq m_k^*$, $\forall s\in S\backslash\mathscr{D}(\epsilon)$. For this we use the following nested induction.

Suppose: (a) $i(s)$ will not deviate from σ^* on $S(k-1)$, $k\geq2$; and, (b) the inequality in (A.2) fails at s, where $k(s)=k$ and $m(s)=m\geq1$. Then that inequality also fails at s', where $k(s')=k$ and $m(s')=m+1$. By (b), waiting to reach $\mathscr{D}(\epsilon)$ from s' is worse for $i(s)$ than waiting for one period, producing for one period, and then following σ^*. While, by (a), producing for one period and then following σ^* is at least as good as producing for one period, waiting for one period, and then following σ^*. Thus, producing and then following σ^* dominates waiting to reach $\mathscr{D}(\epsilon)$ from s'. So either, for all $s\in S(k)$ with $k(s)=k$, producing for one period and then following σ^* is better than waiting to reach $\mathscr{D}(\epsilon)$, whence $m_k^*\equiv0$; or, for some $m_k^*>0$, the inequality in (A.2) holds iff $m(s)\leq m_k^*$. For fixed k, this induction on m implies $i(s)$ follows σ^* on $S(k)$ if it does so on $S(k-1)$, $k\geq2$. The induction on k, and the fact that $i(s)$ follows σ^* on $S(1)$, then imply that the "optimal waiting strategy" defined in (A.2) is firm i's best response on all of S.

We now show that the larger firm will not deviate from constant production on $S(k)$, $k\geq1$, given that the smaller firm follows σ^*. We do so by induction on k. Clearly $j(s)$ will not deviate from σ^* at any node where $i(s)$ is waiting, so we need only consider s for which $m(s)>m_{k(s)}^*$. We argued in the text that $j(s)$ follows σ^* on $S(1)$. Assume this is so on $S(k-1)$, $k\geq2$. Take any $s\in S(k)$ such that $k(s)=k$ and $m(s)>m_k^*$. Let s' be the node reached if $j(s)$ follows σ^*, let s'' be the node reached if $j(s)$ deviates from σ^*, and let $m\equiv m(s')$. There are three cases: (1) $m>m_{k-1}^*$; (2) $m<m_{k-1}^*$; and (3) $m=m_{k-1}^*$.

Case (1): By the definitions, $m(s'')=m+1$. So in this case $i(s)$ will be playing "1" at the successor to node s, whether $j(s)$ follows σ^* or not. But then $j(s)$ cannot gain by deviating. For, given discounting, firm j does better to follow σ^* at s and then deviate at s', than to deviate at s and then follow σ^* at s'', since both paths terminate after two periods at the same node. Yet, by the induction hypothesis, the payoff from the former course is a lower bound on the value of following σ^* at s, and the payoff to the latter course is the best that can happen if firm j deviates at s.

Case (2): In this case $i(s)$ will be waiting to reach $\mathscr{D}(\epsilon)$ from the successor to s, whether $j(s)$ follows σ^* or not. So $j(s)$ does better under σ^* from s'' than it would do under the pure alternating profile. Moreover, since $s\notin\mathscr{D}(\epsilon)$, it follows that:

$$V_{j(s)}(\sigma^*)(s'')>V_{j(s)}(\hat{\sigma})(s'')>\beta(\frac{1-\epsilon}{1-\delta}).$$

So, were $j(s)$ to play "0" at s, its value there would exceed $\delta\beta(1-\epsilon)/(1-\delta)$. But then the

reasoning of Theorem 2 and Corollary 1 implies firm j will not wait to "take its turn" under such circumstances. Hence, j follows σ^* at s.

Case (3): In this case if $j(s)$ follows σ^* it goes to the node s' where $i(s)$ begins waiting to reach $\mathscr{D}(\epsilon)$, while by deviating it goes to s'' where $i(s)$ is not waiting. We know from Case (2) that if $i(s)$ were to begin waiting at s'' then it would not pay for firm j to deviate at s. Thus, we can dispose of this case if we can show that firm j is worse off at s'' when firm i produces there than when firm i begins waiting there. Let t' denote the first date at which firm $i(s)$ waits along the path $\{s_t(s'', \sigma^*)\}$, $0 \leq t' \leq k-1$.

If $t'=0$, we are in Case (2). A critical fact for this proof is the following:

$$m(s_{t'}(s'', \sigma^*)) \leq m(s''), \quad 1 \leq t' < k-1.$$

The waiting time to reach $\mathscr{D}(\epsilon)$ does not rise as one moves along the path $\{s_t(s'', \sigma^*)\}$.

To see this notice that, by Corollary 2, $m(s)=\text{Min}\{m \in I \mid \forall k, 1 \leq k \leq k(s), \phi(n(s)-m+1,k) \geq \epsilon\}$. The discussion following Corollary 2 in the text implies that, as $k(s)$ falls, the critical value of n for which $\phi(n,k) \geq \epsilon$, $1 \leq k \leq k(s)$, does not decline. But $n(s_{t'}(s'', \sigma^*))=n(s'')$, and $k(s_{t'}(s'', \sigma^*))<k(s'')=k-1$, by the definitions. Thus, $m(s_{t'}(s'', \sigma^*)) \leq m(s'')$.

With this observation in hand, we complete the proof by showing that firm j is worse off under σ^* at s'' if firm i produces there, than if it does not. Given that firm i is about to begin waiting for $\mathscr{D}(\epsilon)$, it is obviously better for firm j if the wait is longer. Define $V^*(t) \equiv V_{j(s)}(\sigma^*)(s'')$, given that $t'=t$; let $s_t \equiv s_t(s'', \sigma)$, and $m_t \equiv m(s_t)$. We need to show that $V^*(0) \geq V^*(t)$, $1 \leq t \leq k-1$. From the definition of the profile σ^* we have:

$$V^*(t) = \beta \epsilon \left(\frac{1-\delta^t}{1-\delta}\right) + \beta \delta^t \left(\frac{1-\delta^{1+m_t}}{1-\delta}\right) + \delta^{1+t+m_t} \cdot V_{j(s)}(\hat{\sigma})(\hat{s}_t), \quad 0 \leq t \leq k-1,$$

where $\hat{s}_t \in \mathscr{D}(\epsilon)$ is the first node reached at which firm j waits. Since a longer wait by i favors j, and since $m_{t+1} \leq m_t$, we know that $V^*(t+1)$ is no greater than the value implied above when $m_{t+1}=m_t$ is assumed. But $m_{t+1}=m_t$ implies $\hat{s}_{t+1}=\hat{s}_t-(1,1)$. Therefore:

$$V^*(t)-V^*(t+1) \geq \beta \epsilon \left\{ \left(\frac{1-\delta^t}{1-\delta}\right) - \left(\frac{1-\delta^{t+1}}{1-\delta}\right) \right\} + \beta \left\{ \delta^t \left(\frac{1-\delta^{1+m_t}}{1-\delta}\right) - \delta^{t+1} \left(\frac{1-\delta^{1+m_t}}{1-\delta}\right) \right\}$$
$$+ \delta^{1+t+m_t} \{ V_{j(s)}(\hat{\sigma})(\hat{s}_t) - \delta V_{j(s)}(\hat{\sigma})(\hat{s}_t-(1,1)) \}, \quad 0 \leq t < k-1.$$

Now because $\hat{s}_t \in \mathscr{D}(\epsilon)$, and j is waiting at $\hat{s}_t$, we have: $V_j(\hat{\sigma})(\hat{s}_t) \geq \beta \epsilon + \delta V_j(\hat{\sigma})(\hat{s}_t-(1,1))$. Substituting this inequality into the expression above yields the following:

$$V^*(t)-V^*(t+1) \geq \beta \delta^t (1-\epsilon)(1-\delta^{1+m_t}) > 0, \quad 0 \leq t < k-1.$$

So $V^*(0)>V^*(t)$, $1 \leq t \leq k-1$. We conclude that j will not deviate from σ^* in Case (3). □

Proof of Theorem 4:

(i) We begin by noting that $\bar{N}$ is closed to deviation from $\tilde{\sigma}$, in view of the fact that $\bar{n}_k \geq \bar{n}_{k-1}$, $k \geq 1$. Thus we must show that neither firm can gain by deviating at any node

$s \in \bar{N}$. Neither will deviate when it alone produces. The proof of Theorem 2 implies $j(s)$ will not deviate from waiting at nodes $s \in \bar{N}$ for which $n(s) = \bar{n}_{k(s)} - c_{k(s)}$, since its payoff at the successor node under $\tilde{\sigma}$ does not exceed $\delta \beta (\frac{1-\epsilon}{1-\delta})$. So what must be shown is that, for some $\{\bar{m}_k\}$, $i(s)$ will not deviate when $\tilde{\sigma}_{i(s)}(s) = 0$, and neither firm will deviate when $\tilde{\sigma}(s) = (1,1)$. The structure of the argument is similar to that used for Theorem 3.

Consider first firm $i(s)$. We assert the existence of some sequence of integers $\{\bar{m}_k\} \subset I$, yet to be determined, such that "$\tilde{\sigma}_{i(s)}(s) = 0$ iff $\bar{n}_{k(s)} - c_{k(s)} < n(s) \leq \bar{n}_{k(s)} + \bar{m}_{k(s)}$" is a best response to "$\tilde{\sigma}_{j(s)}(s) = 1$ iff $\bar{n}_{k(s)} - c_{k(s)} < n(s)$," $s \in \bar{N}$. Let $\bar{N}^k \equiv \{s \in \bar{N} \mid k(s) \leq k\}$, $k \in I$. A nested induction of the sort used in the proof of Theorem 3 establishes that the optimal behavior of firm i must be of the "optimal waiting" form: I.e., if for some $s \in \bar{N}^k \backslash \bar{N}^{k-1}$ with $\bar{n}_k - c_k < n(s)$ it is optimal for i to produce, then it is also optimal for i to produce at $s' \in \bar{N}^k \backslash \bar{N}^{k-1}$ for which $n(s') > n(s)$. The argument turns on the fact that, due to discounting, it could never pay i to wait in order to reach a node at which i will produce, when it could produce and reach a node from which it might choose to wait. So we need to show that the length of i's optimal wait on $\bar{N}^k \backslash \bar{N}^{k-1}$ is at least c_k. This is obviously true for $k=1$, since i would wait a total of n' periods to play its one unit, so (consulting (4.7)) $\bar{m}_1 = n' - c_1 > 0$. Suppose, inductively, that $\{\bar{m}_1, ..., \bar{m}_{k-1}\} \subset I$, $k \geq 2$, are such that $\tilde{\sigma}$ gives i's best play on $\bar{N}^{k-1}$, and consider his optimal behavior at $s \in \bar{N}^k \backslash \bar{N}^{k-1}$.

Now waiting by i takes it to a node s'_k for which $n(s'_k) = \bar{n}_k - c_k$, after which $\tilde{c}$ causes the following pattern of play to ensue: i produces for one period, j produces for c_{k-1} periods, i produces for one period, j produces for c_{k-2} periods, etc. We know that $V_j(\tilde{\sigma})(s'_k) = \delta V_{k-1}$, for V_{k-1} given in (4.7). Let $T_k \equiv 2k + \bar{n}_k - c_k = n' + k + (c_1 + ... + c_{k-1})$ be the sum of firms' capacities at s'_k. Therefore firm i's payoff at s'_k is given by:

$$\text{(i)} \quad W_k \equiv V_i(\tilde{\sigma})(s_k') = \beta(1 - \delta^{T_k})/(1-\delta) - \delta V_{k-1}.$$

Assuming, inductively, that $V_i(\tilde{\sigma})(s)$ is well defined on $\bar{N}^{k-1}$, we conclude that waiting will be optimal at $s \in \bar{N}^k \backslash \bar{N}^{k-1}$ if and only if:

$$\text{(ii)} \quad \delta^{m(s)} W_k \geq \beta\epsilon + \delta \cdot V_i(\tilde{\sigma})(s-(1,1)) = \underset{1 \leq t \leq k}{\text{Max}} \{\beta\epsilon(\frac{1-\delta^t}{1-\delta}) + \delta^{m(s)+c_{k-1}+...+c_{k-t}} \cdot W_{k-t}\},$$

where $m(s) \equiv n(s) - \bar{n}_k + c_k$, and $W_0 = c_0 \equiv 0$. The inequality in (ii) states: "i prefers to wait at s." The equality in (ii) uses the induction hypothesis to infer that, conditional on i producing at s, i's optimal behavior on $\bar{N}^{k-1}$ takes the form: continue to produce until, with $k-t$ (possibly zero) units left, it pays to wait for the node s'_{k-t}, and associated payoff W_{k-t}." We need to show that the inequality in (ii) holds when $m(s) = c_k$, so as to

conclude that i waits at $s \in N^k \backslash N^{k-1}$ if and only if $n(s) \leq \bar{n}_k + \bar{m}_k$, for some $\bar{m}_k \in I$.

In the proof of Lemma 3 we noted that $\beta(\frac{1-\epsilon}{1-\delta}) \geq V_k > \beta(\frac{1-\epsilon/\delta}{1-\delta})$, $k \geq 1$. So by (i):

$$\text{(iii)} \quad \beta[1 + \frac{\epsilon - \delta^{T_k}}{1-\delta}] > W_k \geq \beta[1 + \frac{\delta\epsilon - \delta^{T_k}}{1-\delta}], \ k \geq 2.$$

Substituting (iii) into (ii) and rearranging, we see that (ii) will hold if, for $1 \leq t \leq k$:

$$\text{(iv)} \quad \delta^{c_k} \cdot [1 - \delta + \delta\epsilon - \delta^{T_k}] \geq \epsilon(1 - \delta^t) + \delta^{c_k + \ldots + c_{k-t}} \cdot [1 - \delta + \epsilon - \delta^{T_{k-t}}].$$

Recalling the definition of T_k and using a bit of algebra we have that (ii) holds if:

$$\text{(v)} \quad \delta^{c_k} \cdot (1 - \delta + \delta\epsilon) \cdot [1 - \delta^{c_{k-1} + \ldots + c_{k-t}} \cdot (\frac{1 - \delta + \epsilon}{1 - \delta + \delta\epsilon})] \geq [\epsilon - \delta^{T_k + c_{k-t}}] \cdot (1 - \delta^t), \ 1 \leq t \leq k-1.$$

Now by (4.7)(ii), $\delta^{c_k} \cdot [1 - \delta(1 - \delta)V_{k-1}] \geq \epsilon$ and $V_{k-1} \leq (\frac{1-\epsilon}{1-\delta})$, $k \geq 1$; so, $\delta^{c_k}(1 - \delta + \delta\epsilon) \geq \epsilon$, $k \geq 2$. By Lemma 3, $c_k \geq 1$, $k \geq 1$. Obviously, $(\frac{1 - \delta + \epsilon}{1 - \delta + \delta\epsilon}) < \frac{1}{\delta}$. So, if $c_{k-1} + \ldots + c_{k-t} \geq t+1$, then we may conclude that $[\text{LHS of (v)}] > \epsilon \cdot [1 - \delta^t] > [\text{RHS of (v)}]$, and (ii) holds.

Suppose $c_{k-m} = 1$, $1 \leq m \leq t$. Then by (4.8), $\delta^2/(1 + \delta + \delta^2) < \epsilon \leq \delta/(1+\delta)$, and $c_k \leq 2$. Also, $W_k = \beta(1 - \delta^{2t})/(1 - \delta^2) + \delta^{2t}W_{k-t}$, and $W_{k-t} \leq \beta(1 - \delta^{2(k-t)})/(1 - \delta^2)$. So (ii) holds if, for $1 \leq t \leq k$:

$$\delta^{c_k}[\frac{1 + \delta^t}{1 + \delta} - \beta^{-1}\delta^t(1 - \delta)W_{k-t}] \geq \epsilon, \text{ which holds if } \delta^{c_k}[\frac{1 + \delta^{2k-t}}{1 + \delta}] \geq \epsilon, \text{ which holds if } \frac{\delta^2}{1 + \delta} \geq \epsilon.$$

But (4.8) implies that if $\epsilon > \delta^2/(1+\delta)$, then $c_k = 1$, $\forall k \geq 1$. Since $W_k \leq \beta(\frac{1-\epsilon}{1-\delta})$, firm i will wait at least one period to reach an alternating path. Thus we have proven that (ii) holds for $1 \leq m(s) \leq c_k$. The integer $\bar{m}_k$ is given by: $\bar{m}_k \equiv \text{Max}\{m(s) \mid (ii) \text{ holds}\} - c_k \geq 0$. We conclude that i's best play on $\bar{N}$ is given by $\tilde{\sigma}_i$, for $\{\bar{m}_k\}$ inductively defined in this manner.

It remains to show that firm $j(s)$ cannot gain by deviating from $\tilde{\sigma}$ at $s \in \bar{N}$ for which $n(s) > \bar{n}_{k(s)} + \bar{m}_{k(s)}$, and $\tilde{\sigma}(s) = (1,1)$. We sketch the argument here, as the reasoning closely parallels that employed in this part of the proof of Theorem 3. Again, there are three cases, distinguished by the behavior of $\tilde{\sigma}$ at the nodes s' and s'', respectively reached if j does or does not follow $\tilde{\sigma}$ at s. If $\tilde{\sigma}(s') = (1,1)$ then j will not deviate, since that involves waiting to reach a node where it will produce, when it could produce and reach a node from which it might wait. If $\tilde{\sigma}_j(s'') = 0$ then j will not deviate, since by the definition of $\bar{n}_k$, the fact that $\tilde{\sigma}(s) = (1,1)$ implies $V_j(\tilde{\sigma})(s)$ is too large for j to forebear at s in order to reach s'', where i forebears. There remains the case in which $\tilde{\sigma}_i(s') = 0$, but $\tilde{\sigma}(s'') = (1,1)$, so by deviating j avoids reaching a node where i would have begun to wait. Argument analogous to that used for Theorem 3 disposes of this case as well.

The point, as in the proof of Theorem 3, is that j never gains by delaying i's wait. Bear in mind that i is waiting to reach the unique admissible path in $\bar{N}$ at some node s'_k

for which $n(s_k') = \bar{n}_k - \bar{c}_k$. If i were to begin waiting for s_k' from a node $\bar{s}$ for which $k(\bar{s}) = k$ and $n(\bar{s}) = \bar{n}_k - \bar{c}_k + m$, then the payoff to j as a consequence of this waiting by i is: $\beta(1-\delta^m)/(1-\delta) + \delta^{m+1}V_{k-1}$, for V_{k-1} as defined in (4.7). If, however, i were to begin at $\bar{s}-(1,1)$, then the implied payoff to j at $\bar{s}$ is: $\beta\epsilon + \delta\beta(1-\delta^{m-1})/(1-\delta) + \delta^m V_{k-1}$. The latter payoff is less than the former since, by (4.7), $\delta^m V_{k-1} < \beta(\frac{1-\epsilon}{1-\delta})$. Thus, if j deviates from $\tilde{\sigma}$ at s, it would strictly prefer that i begin waiting at s'' than at $s''-t\cdot(1,1)$, for any $t \geq 1$. And yet, it would rather follow $\tilde{\sigma}$ at s than deviate and reach s'' with i beginning to wait there. So j will not deviate from $\tilde{\sigma}$ under any contingency, and we may conclude that $\tilde{\sigma}$ is a restricted MPE on $\bar{N}$.

(ii) We now show that $\{s \in S \mid n(s) \leq \bar{n}_{k(s)} + \overline{m}_{k(s)}\} \subset M^*$. Since $\tilde{\sigma} \in \Sigma^*_{|\bar{N}}$, Theorem 1 implies that there is some $\overline{\sigma} \in \Sigma^*$ such that $\overline{\sigma} \equiv \tilde{\sigma}$ on $\bar{N}$. By definition of the profile $\tilde{\sigma}$ and the sequence $\{\overline{m}_k\}$ we know that $\{s_t(s,\overline{\sigma})\} \subset M(\overline{\sigma})$, for all $s \in \bar{N}$ for which $n(s) \leq \bar{n}_{k(s)} + \overline{m}_{k(s)}$. So we will have proven the Theorem if we can show that, for every $s \in S \backslash \bar{N}$, there is some $\overline{\sigma} \in \Sigma^*$ which generates a collusive path from s. We do so by constructing an extension of $\tilde{\sigma}$ which has the property that it generates an alternating path from s until $\bar{N}$ is reached, after which it follows $\tilde{\sigma}$.

Note that $\{s_t(s,\hat{\sigma})\} \cap \bar{N} \neq \emptyset$, $\forall s \in S \backslash \bar{N}$, since $(a,0) \in \bar{N}$ and $(0,b) \in \bar{N}$, $\forall a,b > 0$. For $s \in S \backslash \bar{N}$, let $\bar{s}$ be the first node reached from s along $\{s_t(s,\hat{\sigma})\}$ such that $\bar{s} \in \bar{N}$. Then it must be that $n(\bar{s}) \leq \bar{n}_{k(\bar{s})}$. If $k(\bar{s}) > 0$ and $\hat{\sigma}(\bar{s}) = \tilde{\sigma}(\bar{s})$ we will say that $\hat{\sigma}$ is "compatible" with $\tilde{\sigma}$ at $\bar{s}$. When $\hat{\sigma}$ and $\tilde{\sigma}$ are not compatible at $\bar{s}$, then redefine $\hat{\sigma}$ so that firm A produces when the sum of stocks is odd, instead of even. With this redefinition in mind we can state: $\forall s \in S \backslash \bar{N}$ there exists an alternating profile $\hat{\sigma}$ which is compatible with $\tilde{\sigma}$ in the sense that they both have the same firm producing at the first node reached from s under $\hat{\sigma}$ that lies in $\bar{N}$. With $\hat{\sigma}$ so defined, let $\mathscr{A}(s) \equiv \{s_t(s,\hat{\sigma})\} \cup \bar{N}$ and let $\sigma[s]$ be a strategy profile defined on $\mathscr{A}(s)$ so that: $\sigma[s](s') = \tilde{\sigma}(s')$, $\forall s' \in \bar{N}$, and $\sigma[s](s') = \hat{\sigma}(s')$, $\forall s' \in \mathscr{A}(s) \backslash \bar{N}$. In words, for given $s \in S \backslash \bar{N}$, $\sigma[s]$ is the profile defined on the union of $\bar{N}$ and the compatible alternating path from s, such that it generates alternation outside of $\bar{N}$, and follows $\tilde{\sigma}$ on $\bar{N}$. It is obvious that $\mathscr{A}(s)$ is closed to deviation from $\sigma[s]$, and that $\sigma[s]$ generates a collusive path from s. We now show that $\sigma[s] \in \Sigma^*_{\mathscr{A}(s)}$.

Clearly we need only be conerned with points $s' \in \mathscr{A}(s) \backslash \bar{N}$. The proof of Theorem 2 showed that it pays to deviate from alternation only when the payoff at the node reached without deviating exceeds $\beta(\frac{1-\epsilon}{1-\delta})$. Neither firm's payoff exceeds this bound at $\bar{s}$. At any

predecessor of $\bar{s}$ and successor of s along the path $\{s_t(s,\hat{\sigma})\}$, each firm's payoff consists of

receiving β on alternate dates until the node $\bar{s}$ is reached, and hence may be written: $\beta(1-\delta^{2t})/(1-\delta^2)+\delta^{2t}V_j(\sigma)(\bar{s})\leq\beta(1-\delta^{2t})/(1-\delta^2)+\delta^{2t}\beta(\frac{1-\epsilon}{1-\delta})\leq\beta(\frac{1-\epsilon}{1-\delta})$, in view of the fact that $\epsilon\leq\delta/(1+\delta)$. So neither firm will deviate from $\sigma[s]$ anywhere on $\mathscr{A}(s)$. Hence $\sigma[s]\epsilon\Sigma^*_{|\mathscr{A}(s)}$. By Theorem 1, $\exists\ \bar{\sigma}\epsilon\Sigma^*$ such that $\sigma[s]\equiv\bar{\sigma}$ on $\mathscr{A}(s)$, and so $s\epsilon M^*$. □

Derivation of Equation (5.2):

We seek $w^*(y)\equiv\lim_{h\to 0}\{h\cdot m^*_{k(s_0(h,R))}\}$. From the implicit definition of $\{m^*_k\}$ given in equation (A.2), and the properties of the MPE σ^* established in the proof of Theorem 3, we conclude that $m(s)=m^*_{k(s)}$ if and only if the following condition holds:

$$(A.3)\quad \beta\delta^{m(s)+1}\cdot\left[\frac{1-\delta^{2k(s)}}{1-\delta^2}\right] \geq \underset{1\leq t\leq k}{\text{Max}}\ \beta\{\epsilon(\frac{1-\delta^t}{1-\delta}) + \delta^t\cdot\delta^{1+m(s-t\cdot(1,1))}\cdot\left[\frac{1-\delta^{2(k(s)-t)}}{1-\delta^2}\right]\}$$

$$> \beta\delta^{m(s)+2}\cdot\left[\frac{1-\delta^{2k(s)}}{1-\delta^2}\right].$$

(A.3) states that firm $i(s)$ does better to wait for $\mathscr{D}(\epsilon)$ from s than to produce (and to continue producing until waiting becomes optimal), but that this would not be so if the wait were one period longer. In view of the fact that $h\cdot k(s_0(h,R))\to y$, as $h\to 0$, we have that $h\cdot m(s)\to w^*(y)$ as $h\cdot k(s)\to y$ in (A.3). In other words, (5.2) is obtained from (A.3) by letting $h\to 0$ in the latter, identifying $w^*(y)$ with $\lim\{h\cdot m(s)\}$, and y with $\lim\{h\cdot k(s)\}$.

The key to the derivation is the observation that $h\cdot t\to 0$ as $h\to 0$. Assuming this is true for the moment, the limit of the RHS of the first inequality in (A.3) can be evaluated using calculus. Introduce the infinitesmals $dy\approx h\cdot t$ and $dw\approx h\cdot[m(s)-m(s-t\cdot(1,1))]$ for $h\approx 0$; let $h\cdot m(s)\to w^*$ and $h\cdot k(s)\to y$ as $h\to 0$. In the text we noted that, by taking limits in (4.4) the waiting time defined in (4.5) satisfies: $h\cdot m(s)\to x+\frac{1}{r}\ln[\frac{1}{2}+(\epsilon-\frac{1}{2})e^{2ry}]\equiv w(x,y)$, as $h\cdot k(s)\to y$ and $h\cdot n(s)\to x$. Therefore, simple differentiation yields the following:

$$(i)\quad dw \approx w(x,y)-w(x,y-dy) \approx \frac{\partial w}{\partial y}(x,y)\cdot dy = \left[\frac{2(2\epsilon-1)e^{2ry}}{1+(2\epsilon-1)e^{2ry}}\right]\cdot dy \equiv [1-u(y)]\cdot dy,$$

where $u(y)\equiv[1-(2\epsilon-1)e^{2ry}]/[1+(2\epsilon-1)e^{2ry}]$. Now, using the definition of $\delta(h)$ and $\beta(h)$ and taking $h\approx 0$, (A.3) becomes:

$$(ii)\quad e^{-rw^*}\cdot[\frac{1-e^{-2ry}}{2r}] \approx \epsilon dy + e^{-rdy}\cdot e^{-r(w^*-dw)}\cdot[\frac{1-e^{-2r(y-dy)}}{2r}].$$

Note that (ii) has the intuitive interpretation that firm i, with capacity y, is indifferent between waiting w^* units of time to share the market, or producing for a short interval and waiting thereafter. Letting $dy\to 0$ in (ii), and using (i) yields the following:

$$(iii)\quad 0 = \epsilon - e^{-rw^*}\cdot[\frac{1}{2}\cdot(1-e^{-2ry})\cdot u(y) + e^{-2ry}].$$

It is easy to see that (iii) implies (5.2).

Notice that u(0)=(1−ϵ)/ϵ<1; u′(y)<0; and u($\overline{y}$)=0, for $\overline{y}$≡−ln(2ϵ−1)/2r. So there

is a unique $\hat{y}$∈(0,$\overline{y}$) satisfying: [2ϵ−u($\hat{y}$)]/[2−u($\hat{y}$)] = e$^{-2r\hat{y}}$. For y>$\hat{y}$ the value of w*(y) implied by (iii) is negative, so the first inequality in (A.3) fails when h is sufficiently small, and hence w*(y)=0. Finally, we must verify the presumption that h·t→0 as h→0, upon which the above argument is based. This presumption is valid if, for h≈0 and firm i indifferent between producing and waiting, producing for any positive interval leaves it at a point where waiting is preferred. This is equivalent to the requirement:

$$\partial[w^*(y)-w(x,y)]/\partial y < 0 \quad \text{whenever} \quad w^*(y)=w(x,y).$$

But by (iii), w*(y) is decreasing, while by (i), w(x,y) is increasing in y. □

References

Dutta, P. and R. Sundaram (1989): "The Tragedy of the Commons? A Complete Characterization of Stationary Equilibria in Dynamic Resource Games," (unpublished, Columbia)

Fershtman, C. and M. Kamien (1987): "Dynamic Duopolistic Competition with Sticky Prices," Econometrica, 55, 1151–1164

Fudenberg, D. and J. Tirole (1983): "Capital as a Commitment: Strategic Investment to Deter Mobility," Journal of Economic Theory, 31, 227–256

Fudenberg, D. and J. Tirole (1986): Dynamic Models of Oligopoly, Harwood Academic Publishers, London

Halperin, A. (1990): "Price Competition and Inflation," (unpublished, MIT)

Hansen, L. P., D. Epple, and W. Roberds (1985): "Linear–Quadratic Duopoly Models of Resource Depletion," Chp. 5 in Energy, Foresight and Strategy, Thomas J. Sargent (ed.), Resources for the Future, Washington, D.C.

Harris, C. and J. Vickers (1985): "Perfect Equilibrium in a Model of a Race," Review of Economic Studies, LII, 193–209

Levhari, D. and L. Mirman (1980): "The Great Fish War: An Example using a Dynamic Cournot–Nash Solution," Bell Journal of Economics 11, 322–334

Loury, G. (1986): "A Theory of 'Oil'igopoly: Cournot Equilibrium in Exhaustible Resource Markets with Fixed Supplies," International Economic Review, 27, 285–301

Maskin, E and J Tirole (1989): "Markov Equilibrium," (unpublished, Harvard)

Reinganum, J. (1981): "Dynamic Games of Innovation," Journal of Economic Theory, 25, 21–46

Reinganum, J. and N. Stokey (1985): "Oligopoly Extraction of a Common Property Natural Resource: The Importance of the Period of Commitment in Dynamic Games," International Economic Review, 26, 161–173

Rosenthal, R. and R. Spady (1989): "Duopoly with Both Ruin and Entry," Canadian Journal of Economics, XXII, 834–851

Simon, L. and M. Stinchcombe (1989): "Extensive Form Games in Continuous Time: Pure Strategies," Econometrica, 57, 1171–1214

Spence, M. (1979): "Investment Strategy and Growth in a New Market," Bell Journal of Economics, 10, 1–19

Sundaram, R. (1989): "Perfect Equilibrium in Non–randomized Strategies in a Class of Symmetric Dynamic Games," Journal of Economic Theory, 47, 153–177

Date Due

Lib-26-67